HOW KNOWLEDGE GROWS

The Evolutionary Development of Scientific Practice

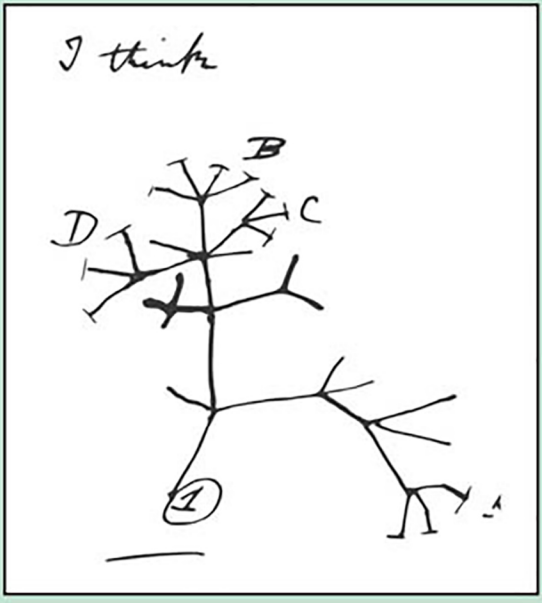

CHRIS HAUFE

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

How Knowledge Grows

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

How Knowledge Grows

The Evolutionary Development of Scientific Practice

Chris Haufe

The MIT Press
Cambridge, Massachusetts
London, England

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

© 2022 Massachusetts Institute of Technology

This work is subject to a Creative Commons CC-BY-NC-ND license. Subject to such license, all rights are reserved.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone by Westchester Publishing Services, Danbury, CT.

Library of Congress Cataloging-in-Publication Data

Names: Haufe, Chris, author.
Title: How knowledge grows : the evolutionary development of scientific practice / Chris Haufe.
Description: Cambridge, Massachusetts : The MIT Press, [2022] | Includes bibliographical references and index.
Identifiers: LCCN 2021057609 | ISBN 9780262544450 (paperback)
Subjects: LCSH: Science—Philosophy. | Science—Methodology. | Knowledge, Theory of.
Classification: LCC Q175 .H35198 2022 | DDC 501—dc23/eng20220405
LC record available at https://lccn.loc.gov/2021057609

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

for Philip Kitcher, veteran

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Contents

Acknowledgments ix

Introduction 1

I The Evolution of Scientific Knowledge 7

1 Again with the Science and the Evolution? 11
2 The Persistence of Practice 45
3 Isolation of the Evolving Population 71
4 Directionality and the Refinement of Scientific Practice 105
5 Branching and the Growth of Scientific Knowledge 141
6 The Recovery of Scientific Knowledge 177

II The Emergence of a Discipline 207

7 The Philosophical Foundations of Disciplinary Autonomy 213
8 Attempted Paleontology 227
9 Wonderful Death 271

Notes 309
References 315
Index 329

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Acknowledgments

This book has profited from exchanges with many different people over many years. I first want to acknowledge David Sepkoski, whose conversation during our time together at the University of Chicago had a formative impact on my thinking about the interaction between social and epistemic forces in the practice of science. David carefully read and discussed part II of the book and encouraged me at each step. In so many ways, this book is the long-digested outcome of his influence, both in person and through his superb book *Rereading the Fossil Record*. I also want to acknowledge the peculiar generosity of John Beatty, who devoted both a week in Vancouver and a week in Cleveland to reading the entire manuscript and giving me the kind of intellectually synoptic feedback that only he can give. John encouraged me to push forward with the project when I considered it too exotic to continue working on. I truly have him to blame for its completion.

Since 2012, I have had numerous discussions about Thomas Kuhn's views, their reception, and their plausibility with (mostly) like-minded philosophers of science who thought the time had come for a fresh look at *The Structure of Scientific Revolutions*. This group includes Paul Hoyningen-Huene, George Reisch, Eric Scerri, Dana Tulodziecki, and Brad Wray. I especially want to thank John Huss for providing invaluable interlocution on everything related to Kuhn and the philosophy of science. John's willingness to chat provided me with a venue in which to think out loud without fear of reprisal, and with the kind of deep knowledge of Kuhn's work that could point me in fruitful directions. I'd also like to single out Matthew Slater, who, while being somewhat less interested than myself in the niceties of Kuhn exegesis, provided significant and singularly insightful feedback at each stage of this project.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

This book also bears the imprint of a number of scholars who indulged me with meetings and responses to emails concerning everything from aspects of scientific practice to bits of theory to who knows what else. Robert Brandon stands alone (in a manner best left unspecified) as an interpreter and articulator of the evolutionary process and the epistemology of evolutionary biology. His conversation during and beyond my year at Duke's Center for Philosophy of Biology fundamentally shaped my perspective on the nature of evolution. Doug Erwin and Dan McShea read part II and fielded hours of questions in person and by email about paleobiology, no branch of which has escaped their mastery. Carl Simpson can explain anything, and I made use of his unique gift over several days in the Smithsonian when he probably should have been doing actual work. Allen Orr read and validated my use of his work on adaptive walks in chapter 4. Peter Godfrey-Smith provided helpful discussion and clarification of his views about Darwinian populations (as did Elliott Sober), without which this book could not have been written. This is particularly heroic, given (as I later discovered) that he considered analogies between science and evolution to be kind of pointless. I hope I have given him reason to think otherwise.

Since 2011, Case Western Reserve University has given me an ideal environment in which to conduct research in the history and philosophy of science. I'd like to gratefully acknowledge the support of all of my colleagues in the Department of Philosophy: Jeremy Bendik-Keymer, Shannon French, Laura Hengehold, Tony Jack, Chin-Tai Kim, and Colin McLarty. Maggie Kaminski and Peter Knox at the Baker-Nord Center for the Humanities provided financial support at several stages. Our historians of science, Alan Rocke and Aviva Rothman in the Department of History, have been more than tolerant respondents as I probed them with questions that generally only philosophers are rude enough to ask. I have been lucky to have many conversations on the nature of disciplines with Michael Clune, a consummate intellectual. Cyrus Taylor in the Department of Physics has been a welcoming colleague and friend, and a valuable source of insight into the culture of physics.

As Dean of the College of Arts and Sciences, Cyrus also facilitated research during the early stages of this project through two W. P. Jones grants, as well as a year away at the University of Rochester Humanities Center, where much of this book was written. I am grateful to the Center's director, Joanie Rubin, for giving me the opportunity to work in the perfect

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

intellectual environment she has manufactured. I was fortunate to have met and conversed with Chris Rovee and Will Miller during my time there, both of whom offered perspectives on their own disciplines that shaped the views for which I argue in this book. I've also benefited from the institutional support of Insight and its CEO and founder, Jawad Shah. Jawad is a truly great friend and an unflagging advocate for my research. His own research experience, critical mind, and insatiable intellectual curiosity have had a significant influence on how I understand the development of scientific knowledge.

Long have I waited to publicly express my gratitude to Marc Lange and Philip Kitcher, who have been continuous sources of wise council and encouragement as I attempted to shape this book into something of philosophical significance. I owe them more than I can ever express or repay. I pounced on Marc after arriving in North Carolina in 2006, and have continued to insert myself into his orbit ever since. His generosity with his time and intellect has had a huge influence on how I approach the philosophy of science. Philip was extremely patient with me as I bounced around the landscape of contemporary philosophy as a young graduate student. That patience gave me the time to mature intellectually. He has never refused a request, no matter how little time I've left him to fulfill it. (He even advised me on the title of this book.) And he has never once made me feel like he's doing me a favor, even though he has done me countless favors. I know from the testimony of several people—many of whom are not even his students—that this is not unusual for him. As I try to imagine the veritable hellscape that is Philip's email inbox, I feel particularly honored by his tireless feedback and support.

I owe a considerable debt to several anonymous referees, each of whom provided carefully considered and detailed feedback. The seriousness with which they took this assignment was exemplary. This book has been fundamentally transformed (for the better) through their efforts to put the brakes on hasty assertions, to point me in helpful directions, and to connect my discussions with adjacent literature. I'm especially grateful to Phil Laughlin, the editor at MIT Press who oversaw the book's production and gave nearly instantaneous feedback on any substantive or editorial questions I had. John Donohue and his editorial team at Westchester Publishing services made wise, careful, and detailed amendments and suggestions during the copyediting stage. Noted Hollywood editor Noor Haydar assisted with the images.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

And finally, I want to acknowledge my family. My deepest gratitude goes to my mother, Merrie, who supported and encouraged my itinerant search for new research opportunities wherever they arose; and to my wife, Maysan, and our children, Zayd, Sinan, and Layla. Raising children has given me many opportunities to reflect on how knowledge is acquired and how it develops, and to use them as (mostly compliant) test subjects outside the reach of any institutional review board as I experiment with different styles of intellectual growth. My wife politely nodded on each of the many occasions where I rushed up to her to excitedly describe yet another analogy between science and evolution. Her historian's sensibilities have been a crucial antidote to the philosopher's tendency to generalize irresponsibly. She has supplied me with alternating doses of encouragement and humility as needed. Her companionship is the most important thing in the world to me. I do not deserve her.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Introduction

Is it possible to say anything general about the nature of scientific inquiry? Thomas S. Kuhn thought so. Sixty years ago, he argued in *The Structure of Scientific Revolutions* that there were three properties exhibited by intellectual communities that achieve consensus. The first is that these communities coalesce around their members' embrace of an intellectual achievement of unprecedented magnitude, one that members subsequently use as a framework for guiding the smaller-scale research questions and methods with which they concern themselves. The second feature of such communities is that their use of this framework takes on a normative character, in that it encourages them to discriminate against ideas—ideas, say, about nature, or about how to investigate it—that differ significantly from those that are widely accepted by community members. The third is that these communities tend to isolate themselves intellectually from other fields of inquiry and from broader social concerns. The possession of these properties, he argued, tends to produce a characteristic historical pattern displayed by the development of scientific knowledge.

How general are these properties across the different branches of modern science? Kuhn believed that they were in some sense (probably for historical reasons) idiosyncratic to the natural sciences, but that one can find these properties instantiated by premodern intellectual communities—"schools" of philosophy, theology, and so on (Kuhn 1962b, 166). Commentators have alleged that his general image of science was essentially extracted from the tradition of physics in which he was trained (Hacking 2012); or maybe even just the period of physics leading up to World War II (Galison 2016). When we look at the many other fields of inquiry that clearly qualify as science, they often seem to bear little resemblance to the image of science by which Kuhn came to be possessed.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Even though the style of scientific inquiry that he described might have been peculiar to prewar physics, I still think there is an important generality lurking here, one that Kuhn understood clearly but that I do not think has been given sufficient attention. What Kuhn seems to have grasped was that the development of intellectual communities often exhibits a certain historical pattern, and that a community's propensity to exhibit that pattern is a function of the degree to which it possesses the three properties mentioned above. Intellectual communities like prewar physics and classical theology—communities that exemplify what it is like for a group to possess those properties to a high degree—fit that historical pattern very well indeed. Many scientific communities do not possess these properties in abundance. Consequently, their development tends not to proceed in quite the manner Kuhn describes. The relevant generalization, then, is one that holds between the degree to which an intellectual community possesses these properties and its propensity to develop in a certain way.

The interesting thing about this suite of properties is that it is precisely the one that determines a biological population's degree of susceptibility to modification by natural selection. Populations that possess in high degrees biological versions of the properties identified by Kuhn tend to be strongly influenced by adaptive modification through natural selection; those that do not tend to be subject to a variety of influences that do not necessarily improve a population's fit with its environment. Equally interesting is the fact that populations under the influence of natural selection tend to exhibit a characteristic pattern of development—a pattern that, surprisingly, closely mirrors the specific course of scientific development that interested Kuhn. Just as there are obviously significant differences across the various sciences along the dimensions he describes, different populations display differences in the kinds of variation they produce, in their exposure to forces of change that do not depend on how well designed they are for their environment, and in the degree to which offspring resemble their parents. And each of these differences contributes to differences in natural selection's influence over them.

I believe that this correspondence is not accidental. Furthermore, I believe that it helps to explain something we take to be epistemically special about scientific knowledge: its tendency to grow both in depth and in breadth. The sciences, along with other cultural endeavors, differ markedly with respect to their propensity to refine existing elements of practice, as

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

well is in their propensity to generate new specializations, which themselves are then subject to a sustained process of refinement. These disparate propensities can be explained in terms of differences in various communities' susceptibility to modification through natural selection—differences that can, in fact, be understood as a function of the differing degrees to which those communities possess the properties that Kuhn regarded as of particular importance to a certain historical pattern of scientific development. There is something epistemically special about this pattern. Our perception that some branches of science are epistemically special is, I believe, partly informed by their instantiation of it.

This correspondence can also contribute to our understanding of why it is that some varieties of human cultural expression are sometimes able to have epistemically distinctive effects despite being the result of the efforts of epistemically compromised agents. Because social phenomena play such a pivotal role in the formation, stability, and trajectory of scientific disciplines, the disciplinary growth of knowledge has frequently been analyzed using the same principles to which we appeal in our attempts to explain other species of human social activity. To this we can add a number of general considerations that make a sociocultural approach the obvious choice for how to proceed in our efforts to understand the historical development of scientific knowledge. These include (1) that science is carried out by self-interested, fallible human beings; (2) that it is a social enterprise; and (3) that the enterprise of science is itself socially embedded in a broader cultural milieu. In addition, science is, especially in our time, powerfully positioned in society in a way that gives it a potent and widely distributed influence over our lives, significantly raising the stakes on the trajectory of scientific activity. Given that scientific investigation must necessarily be afflicted by our cognitive and moral failings, and given that it must be conducted amid all the variegated factors impinging on human life, one might naturally reject as hopelessly misguided any attempt to understand the growth of knowledge that fails to give due credit to the many cognitive impurities that have contributed to its growth. For science is, in Steven Shapin's words, "never pure" (Shapin 2010).

What is the presumed epistemic significance of the fact that science is "never pure"? More specifically, what is the fact that scientific inquiry is a product of "real people"—agents motivated by many sources, many of them nonepistemic, living in a specific place and time, and thus, subject to

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

the modes of thought dominant in that place and time—supposed to tell us about the status and quality of scientific knowledge? The assumption behind this enormously popular approach seems to be that scientific knowledge cannot be epistemically special if its composition is sensitive to noncognitive, individualistic, and cultural influences; those are the wrong kinds of influences for a species of knowledge to have if it is claiming some kind of privileged epistemic status. And given that scientists are just as susceptible to those sorts of influences as the rest of us, scientific knowledge cannot avoid being tainted by the same epistemic impurities as other nonscientific species of knowledge, none of which we regard as epistemically distinguished in any way. There is no refuge from the contaminating effects of time, place, and human weakness. As long as scientific knowledge is "produced by people with bodies, situated in time, space, culture, and society, and struggling for credibility and authority," there can be nothing epistemically special about scientific knowledge.[1]

This critical approach to the epistemic status of scientific knowledge is an instantiation of an important epistemological principle that holds that the epistemic warrant for a belief is undermined when that belief is causally influenced by factors that we all agree should be irrelevant. A nice, albeit terrifying, illustration of this is the current trend of studies that examine affects of various irrelevancies on judges' decision-making behavior. One recent example purports to show that judges impose longer sentences on juvenile defendants after the Louisiana State University (LSU) football team loses a game that they were expected to win, and that the impact is stronger when judges have a bachelor's degree from LSU or when LSU had been ranked in the top 10 (Eren and Mocan 2018). Yikes. However strongly one might feel about LSU football, I'm sure we can all agree that the team's fortunes on any given Saturday ought not to make a difference in how long a child spends in juvenile detention (excluding the unlikely case that the child's act of delinquency caused LSU's loss). If LSU's performance does affect how judges weigh the evidence, then there is something wrong with how that evidence gets evaluated; sentencing would be better, we think, if football was not part of the equation.

When beliefs are affected by factors that are irrelevant to the well-foundedness of inferences, we tend to view those beliefs as epistemically tainted in some way—"impure," in Shapin's terms. That perception of impurity seems to increase with the influential factor's degree of irrelevance.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Football is highly irrelevant to juvenile sentencing; where one grows up is highly irrelevant to whether one's religious commitments are epistemically well founded. And herein lies the epistemic significance of the assertion that science is never pure. What if we were to discover that physicists' interpretation of climate data was just as sensitive to the LSU football team's losses as juvenile sentencing seems to be? What if we were to discover that the propensity for climatologists to agree with each other was as strongly influenced by country of origin as religious belief apparently is? Would this not cast a pall over the epistemic status of scientific knowledge, just as it does over these other phenomena? But this is just how people are. They are influenced by factors that are irrelevant to epistemic well-foundedness. The practitioners of natural science are not immune to such factors. How likely is it that R. A. Fisher would have been a devout Anglican were he to have been born and raised in Mexico? Not bloody likely!

And yet, scientific knowledge grows ever deeper and ever broader, particularly over the past 400 years. As undeniable as the fact that science is done by fallible, selfish, suggestible human beings, the historical development of scientific knowledge has trended unperturbedly toward an increasingly accurate picture of an increasingly large number of phenomena. This trend has persisted not only across many different "people with bodies," not only across differences in "time, space, culture, and society," but even across radically different conceptions of what nature is like and how it ought to be studied. The stability of the growth of scientific knowledge across so much time and so many different cultural hosts has fueled the nagging sense that there is something about modern scientific inquiry itself that makes epistemic growth inevitable. If scientific knowledge has grown despite vast differences among the scientists who contribute to it, perhaps that is because the ability of modern scientific inquiry to deepen and expand our knowledge is often just plain insensitive to variation at the level of individual scientists. Science may never be pure, but the implications of this insight are far from obvious. Often, science seems to be pure enough.

This book is one long argument for the thesis that scientific knowledge persistently grows, even across generations of highly variable groups of scientists, because the development of scientific knowledge is governed by the Darwinian process of descent with modification. Notwithstanding their many differences and failings, scientists—solely by virtue of their efforts to participate in the growth of knowledge—often form groups whose

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

characteristics are modified in response to certain pressures and as a result of certain demographic properties. Their individual attempts to accommodate those pressures result in the unmitigated increases in breadth and depth that we have come to expect from the development of scientific knowledge.

The book's secondary ambition is to support the philosophical claim that this susceptibility to modification through natural selection is one of the things that explains the distinctive epistemic power of certain branches of modern science. If, as Kuhn averred, the peculiar historical pattern of growth displayed by some branches of science is one of the reasons that pattern should be central to our understanding of knowledge itself (Kuhn 1962b, 9), and if that pattern of growth is generated by certain branches' high degree of susceptibility to selection, then the foundation for our views about the epistemic power of science is partly grounded in the fact that its historical development is governed by the process of selection. Had those branches of modern science not been susceptible to Darwinian modification, they would not show that characteristic growth pattern, and consequently, we would not have held them in such esteem. Or so I argue in this book.

We know that scientific knowledge often develops in this way, not because individual scientists are specially endowed with unique truth-finding abilities, but because the properties that expose a group to Darwinian forces are in fact defining features of certain scientific communities. Moreover, the relevant features of these communities are precisely what historians and sociologists of science have been pleading with philosophers of science to recognize for decades as the sine qua non of knowledge production. Scientific communities possess certain cultural norms to which they hold their members. These communities exert a strong influence over the training and professionalization of future initiates. They engage in boundary policing. They reward members for contributing to scientific knowledge in community-validated ways and very rarely otherwise, and they encourage members to use these rewards as motivators. These community-level properties, idiosyncratic though they may be, turn out to be exactly the sort of properties that result in the community's cross-generational stability and in the propagation and adaptive modification of practices within the community from one generation to the next.[2]

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

I The Evolution of Scientific Knowledge

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

To proceed, I begin in chapter 1 by clarifying the explanatory roles that I intend for the Darwinian theory of evolution and by spelling out the burden of proof that will need to be met in order to support the claim that the historical development of scientific knowledge is governed by the Darwinian process. Chapters 2–5 are each devoted to meeting a specific component of that burden. Chapter 6 examines non-Darwinian aspects of scientific change, arguing that the epistemic significance of these aspects can be most clearly articulated against the backdrop of the Darwinian framework developed in previous chapters. Part II (chapters 7–9) is an extended historical case study that carefully illustrates how the branching process described in chapter 5 can explain the emergence of evolutionary paleontology as an autonomous scientific discipline—a new scientific Darwinian population.

There is a delicate art to this sort of argumentation. One can easily slide from a faithful articulation of current evolutionary theory or of scientific practice into an opportunistically vague characterization that facilitates one's argumentative goals. In addition, one has to make stylistic choices regarding how to faithfully represent these goals, choices that can end up having significant downstream effects on the ability to argue for philosophical claims about science in a way that is both recognizable and credible to philosophers, biologists, and the vast panoply of intellectuals interested in the nature of scientific knowledge. I do not claim that my presentation of either evolutionary theory or of the nature of scientific communities has been everywhere optimal. In addition, there are many aspects of both the practice of science and the evolutionary process that I have had to pass over for the sake of other objectives. I hope that I have provided a framework through which more knowledgeable and more talented scholars can profitably explore these issues.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

1 Again with the Science and the Evolution?

> Though the analogy between science and Darwinian evolution is something that people keep coming back to, the analogy has not yielded a lot of new insights so far. We find the same result in many other attempts to describe cultural change as an evolutionary system; a wide variety of processes can be described in a way that borrows from evolutionary biology, but usually this exercise does not teach us anything about those cultural processes that we did not know before. Biological populations have special features that make the abstract concepts of evolutionary theory helpful in trying to understand them. Other systems, which lack these features, can be described in evolutionary terms with a bit of shoehorning, but we do not seem to gain much from doing so.
>
> —Peter Godfrey-Smith, *Theory and Reality* (2003, 166)

1.1 Godfrey-Smith's Gauntlet

Challenge accepted. Why do philosophers of science keep coming back to drink from the Darwinian well? Basically, I think because the analogy between science and Darwinian evolution is, as Kuhn said, "nearly perfect" (1962b, 172). With allowances for "a bit of shoehorning," Godfrey-Smith is likely to agree: both the historical pattern of scientific change as well as the factors that are known to drive that change seem strongly to suggest that the growth of scientific knowledge is governed by the evolutionary process, particularly those aspects of the evolutionary process that were emphasized by Darwin. Why is that not enough?

I want to begin by trying to understand this complaint in a little more detail, if for no other reason than to convince the reader that the philosophical journey on which you are about to embark is not a complete waste of time. The first thing to note is Godfrey-Smith's demand that describing a

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

cultural process "in a way that borrows from evolutionary biology" should teach us something about that cultural process that we did not know before. Before granting the reasonableness of this demand, it's worth pointing out that the demand only makes sense from the perspective of someone specifically interested in gaining an understanding of that cultural process. If, say, we wanted to know about the truly awesome scope of natural selection's influence—as, for example, Dennett (1995) did—then I daresay we would learn something quite profound were we to learn that its influence extends even to the rather unbiological realm of scientific ideas. Who would have thought that the mechanism Darwin devised to explain biological adaptation could also be used to organize our understanding of the growth of knowledge (or the relative growth of wealth [Orr 2018], etc.)?

But if one's goal is, as ours is, to derive insight into the nature of the scientific process, then Godfrey-Smith's demand is not unreasonable, nor is his complaint that the analogy between evolution and science has been relatively fruitless in this regard. Understanding why this demand is reasonable, and why it has not been satisfied by previous attempts, sheds significant light on the project that lies before us. In the context of inquiry, be it philosophical, scientific, or other, the value of an analogy lies in whether and how it enables further understanding—and analogies are not created equal in this respect. An analogy is an offer to organize a certain domain of ideas by using a structure that is native to a different domain (Goodman 1968; Camp 2019). A really good analogy, like Darwin's analogy between selective breeding and natural selection, will provide a structure for thinking about a topic that is intuitively appealing and will make the significance of certain questions obvious and relatively easy to articulate. Many of Darwin's readers had, or had access to, extensive knowledge of the "science" of breeding, and that connoisseurship translated into specific leading questions that allowed readers to explore the extent to which selective breeding actually modeled a mechanism that was operative in nature. Darwin himself anticipates many of these questions in the early chapters of the *Origin*, such as when he discusses whether the structural resemblance between selective breeding and natural selection extends to the phenomenon of *reversion*, whereby a domestic breed will "revert" to resemble their ancestral stock if mating is not properly restricted. Or whether selective breeding is like selection in nature in its capacity to cause any inherited variation to become widespread. Suggestively, another part of the analogy's strength lies

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

in its ability to highlight potential *mis*matches, such as whether the difference between breeder's *intentional* action and the *purely physical* operation of nature is meaningful enough to constitute a difference in mechanism. Indeed, much of the extraordinary rhetorical power of the *Origin*'s argument for natural selection's causal efficacy derives from Darwin's systematic and exhaustive exploration of the sorts of questions that would have naturally occurred to readers concerned with seeing how well selective breeding held up as a model for the transmutation of species in the wild (spoiler alert: it holds up very well).

Let us understand Godfrey-Smith to be making a similar sort of demand on efforts to draw an analogy between the evolutionary process and the process of scientific inquiry. Reframing the scientific process using ideas borrowed from evolutionary biology is of no particular value for understanding science if that new frame does not lead to insights into the process of science to which we might not have otherwise been privy. It need not lead directly to deeper understanding—that is, it is not required that redescribing the scientific process in evolutionary terms by itself solve the philosophical problems associated with the production of scientific knowledge. But it must provide us at least with a sense of what some of the significant philosophical problems are, and of roughly how to approach those problems.

And so, to that extent, there is a certain justice in Godfrey-Smith's criticism: previous attempts to pursue this analogy have failed to deliver on this minimal demand, despite the fact that—at least in my opinion—the resultant redescriptions have been reasonably accurate. There is more to a fruitful analogy than mere accuracy, and understanding what has prevented previous versions of the analogy from bearing fruit will give us some purchase on what form the analogy might need to take in order to repay an investment of readers' time. Before engaging in the time-honored ritual of "the Pummeling of the Predecessors," though, let me make an observation that bears in a general way on the issue of why prior iterations of the analogy have been somewhat lackluster. Plainly put, prior attempts did not have the benefit of the past few generations to help frame their evolutionary approach to scientific knowledge. Each of these attempts, which actually vary quite a lot in their degree of systematicity and comprehensiveness, took place prior to three key developments that bear directly on our subject. The first concerns the substantial gains in clarity, detail, structure, and scope associated with central concepts in evolutionary biology that

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

occurred in the philosophy of biology over the past four decades. I will have more to say near the end of this chapter on specific features of the concept of a Darwinian population that make it a particularly advantageous host for the analogy between science and evolution. Suffice it for now to say that, as a result of being on the other end of so much pathbreaking philosophical research, we are now in a better position than our predecessors to make productive use of the analogy.

The second key development has been the general acceptance of the idea that social or communal features of scientific practice make an enormous difference to the direction and character of scientific research—an acknowledgment that is captured by the slogan "science is a social process." I regard the introduction of sociocultural considerations into the study of scientific knowledge as the most important advance seen during the twentieth century with respect to our general picture of how scientific knowledge develops. More relevant to present purposes, much of the evolutionary analogy's ability to shed light on the philosophical problems surrounding scientific knowledge depends on recognizing the significant influence of sociocultural factors in science. Although analogizers of the past were aware of these developments, and although each makes use of sociocultural considerations in varying ways, these considerations had not yet become embedded in the intellectual fabric of philosophy of science as they are today, and our understanding of the character of that sociocultural influence had not really entered a state of steady maturation until the 1980s—in other words, after the most notable attempts to connect science and evolution had been made. The one exception is David Hull, whose use of sociocultural considerations is bizarrely sophisticated, given the time at which it was written (but not given its author).

The third development concerns the major expansion in our understanding of the evolutionary process that has occurred over the past 40 years in particular. The evolutionary phenomena on which my account leans most heavily—ecological speciation, niche construction, and mass extinction—were not even part of evolutionary theory in the 1970s, let alone the essential components that they are today. Our general understanding of evolutionary diversification is of a still more recent vintage. The thinkers whose work I will canvas in this chapter were drawing on a much more restricted set of evolutionary ideas, and the relatively limited explanatory scope of their accounts reflects this.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

In summary, the significant progress made over the past four decades in our understanding both of evolution and of the nature of scientific inquiry gives us reason to expect that renewed efforts to draw an analogy between science and evolution might be able to teach us far more about science than the attempts with which Godfrey-Smith was unimpressed. Progress in evolutionary theory has clarified and expanded the kinds of phenomena for which scientific analogues might be sought, while progress in our thinking about the nature of science has enriched the set of candidate features that might underlie the ability of science to function as an evolutionary system in a fairly straightforward sense.

One final general observation worth making is that none of the previous attempts in this arena employs evolutionary theory as a framework for *solving philosophical puzzles* associated with scientific knowledge. Nothing about science is depicted as strange or in need of explanation. There are two principle epistemic phenomena with which these attempts occupy themselves: scientific problem-solving (Karl Popper, Donald Campbell, Stephen Toulmin), and conceptual change (Kuhn, Toulmin, Hull). These are very interesting, central features of science, but no effort is ever made to motivate them as phenomena that are *in need of explanation*. Rather, the whole goal seems to be just to show that these "processes can be described in a way that borrows from evolutionary biology . . . with a bit of shoehorning." No wonder "we do not seem to gain much by doing so." Perhaps this goes more directly to Godfrey-Smith's complaint about the fruitlessness of the evolutionary exercise than the relatively inopportune timing of earlier efforts.

A big part of what I hope to accomplish in this book is to show that well-known features of scientific knowledge and the historical pattern carved out by its development—for example, problem-solving, progressive scientific change—should strike us as highly surprising, and that we should *expect their absence unless science is driven by something like the Darwinian process*. Indeed, as many historians and sociologists of science are wont to argue, these putative aspects of the scientific process *are* absent—or, at least, these aspects lack the epistemic character that scientists and philosophers often attribute to them. Although I do not find the arguments for their absence persuasive, I do think that the thesis has enough plausibility to be taken seriously as a framework for interpreting trends in scientific practice. There is an explanatory burden on philosophers of science to show why,

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

despite being the product of the same species that engages in witchcraft, bowling, TikTok, and violent outbursts of jealous rage, science somehow manages to rise above all of our irrational, selfish, destructive tendencies to produce increasingly superior answers to an ever-expanding set of increasingly difficult questions. Why is science—and only science—able to pull off this feat? Ultimately, I argue that it is able to avoid the fate of other human endeavors precisely because science often functions as a Darwinian population. The differences between science and other cultural processes are as real as they appear, I argue, because the sciences tend to more closely approximate Darwinian populations than do other groups.

1.2 Evolutionary Predecessors

I now want to look in more detail at earlier work on the analogy between science and evolution, emphasizing where they fall victim to Godfrey-Smith's complaint and why I swear that things are going to be different this time. Aside from the general misfortunes of historical circumstance discussed in the previous section, each of these approaches has certain deficits that can function as an example of something that a successful version of the evolutionary analogy will have to avoid. Of course, most analogies suffer from some imperfection or other; as Renzi and Napolitano (2011) have shown, the analogy between science and evolution is beset with perhaps more than its share of mismatches. But the criterion of success for an analogy should be how much we gain from it. And it may turn out that even a significant number of analogical incongruities do nothing to impede the production of certain fruit that an analogy is apt to bear. The productive use of analogies in the history of science, for instance, has never been dependent on a perfect structural match between source and target; productive analogies are typically far from perfect. The power of such analogies has always resided in what they offer researchers by way of organizing a domain of phenomena, generating new questions, and suggesting how to begin answering them. This effect can be obtained even with the analogical relationship is relatively coarse grained or limited in scope. Accordingly, I try to focus the discussion below on what in particular has prevented previous attempts from being the kind of generative analogy that Godfrey-Smith rightly demands. I believe that, in the most telling cases, the problems with previous accounts can be traced back to one of the key developments I mentioned above.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

1.2.1 Thomas Kuhn

I view my efforts in this book largely as an attempt to organize and clarify Thomas Kuhn's ideas about scientific knowledge under the evolutionary framework that he ultimately came to see as responsible for the epistemic phenomena that interested him. Although Kuhn never made a systematic effort to explore the evolutionary analogy, the central themes of his *Structure of Scientific Revolutions* scintillate with evolutionary connotations. A sequence of careful readings of that book combined with a career-long fascination with evolutionary theory convinced me that something like an evolutionary version of *Structure* could be made to work. Much of the motivation for the present work stems from an interest in seeing how closely the process described by Kuhn can be modeled as an evolutionary process.

Although Kuhn mentions the analogy between evolution and the process he describes in *Structure*, the idea of science as an evolutionary system is something that he would come to see as centrally important only after his revolutionary account was published. His later work is thus variously peppered with brief evolutionary glosses of a few aspects of scientific practice. Mostly he seems to have been concerned with how evolutionary notions—in particular, those associated with speciation—could make sense of incommensurability. In this regard, I think that the history of Kuhn's ideas is, in a way, pretty tragic. I have tried to make clear throughout this book that, as a causal model, I take Kuhn's account of the development of scientific knowledge very seriously. So do lots of other people; on average, it has been cited nearly once a day since it was published (Abbott 2016, 173). No other single account captures as many phenomena as Kuhn's, and most alternative accounts of any one of the phenomena he captures are nowhere near as elegant. After more than 50 years, I believe Kuhn's original model remains one of the most valuable products of twentieth-century humanistic thought. And I would bet that another book as rich and as elegant as *Structure* will not be produced in my lifetime.

The tragedy—for me, anyway—is that this is not the book's legacy. And the reason it is not the book's legacy is mostly because of one little idea: *incommensurability*. Kuhn was palpably enamored of incommensurability when he wrote *Structure*, and he devoted a considerable amount of space to spelling the idea out and to showing why it makes revolutions necessary. For all that, though, incommensurability plays a relatively minor role in the overall causal model presented in the book. It briefly drops in to explain

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

revolutionary change; the rest of the work is done by his notion of paradigm-governed normal science.

And yet the book's reception was overwhelmingly dominated by philosophers' negative reactions to incommensurability. Their fundamental fear was that the incommensurability of paradigms would imply that a choice among them could not be rationally justified. Therefore, on Kuhn's account, a major and routine feature of the historical development of science was not rationally justifiable. Therefore, what had been held up as a historical sequence of epistemic improvements was simply groups of dogmatic practitioners trading in one arbitrary set of commitments for another arbitrary set whose superiority they could—by necessity—neither justify nor even articulate. This was not to be tolerated, especially in a sociocultural moment (the 1960s) in which the authority of many cherished institutions was increasingly under threat. To add insult to injury, as George Reisch has persuasively argued, it certainly did not help that normal science—the only clearly progressive part of Kuhn's model—literally depended on practitioners' blind faith rather than their unrelenting open-mindedness (Reisch 2005). Kuhn turned our image of science inside out, and he did it with such alarming elegance that no amount of effort could be spared in undermining his story.

In response, Kuhn would spend much of the post-*Structure* period of his career attempting to systematize the phenomena that were symptomatic of incommensurability, hoping to clarify to others as well as to himself precisely what the idea was and why it appeared to him to be of fundamental importance for understanding the nature of scientific inquiry. Through this process of refinement, he would eventually abandon the description of revolutions as "paradigm shifts." In its place, he developed a much narrower conception of revolutionary change as lexical or taxonomic change—that is, as the replacement of one conceptual ordering of nature with another.

This final conception preserved many of the features that Kuhn found valuable about his original framework of paradigm shifts while significantly narrowing the scope of the kind of change that characterized a scientific revolution. First, it allowed Kuhn to retain the psychological or perceptual aspect of revolutionary change that grounded his attachment to the idea that practitioners with different taxonomies are "living in different worlds"; it would be a mistake to underestimate Kuhn's enduring esteem for the explanatory power of this thesis.[1] At the same time, it explained how this psychological phenomenon could result in a group-level dispute

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

over how science is to be done. In the years after the publication of *Structure*, Kuhn's views on how practitioners acquire their paradigms came to rely more and more heavily on the general process of language acquisition. In particular, he began to see the training of a scientist as, more than anything, the process of learning a very specialized lexicon to be employed in the investigation of nature.

The fact that Kuhn never abandoned incommensurability even in the face of very widespread, cogent, and persistent criticism is, at least to me, deeply significant. Kuhn was a smart man, and he was a very good, very patient reader when it came to criticism (which is in itself remarkable, considering how much he received). He took a lot of care in responding to critics. His responses typically reflect a deep understanding of the complaint, and he more often than not accepts responsibility. But his approach, at least when it came to incommensurability, was always to look for some way to refine or better express the idea that had come under dispute. It seemed that, as long as there were serious critics who doubted the incommensurability thesis, he was never going to drop it. It is almost painful to try to imagine all the different sorts of interesting big ideas Kuhn's three highly productive decades after *Structure* might have generated had the opposition to incommensurability been less pronounced. In no other dimension is this effect more detectable, I think, than in his failure to develop systematically the evolutionary analogy.

I have tried to show that much of what Kuhn had to say about the maturation and structure of scientific inquiry—the penetrating, synoptic image of science that is now very widely accepted—follows directly from conceiving of scientific practice as a paradigmatically Darwinian population; this includes some version of incommensurability. This was in no way an attempt to honor Kuhn. The point has rather been to show that the evolutionary model of scientific inquiry implies the existence of a wide range of phenomena that most of us now take for granted. This is probably the most compelling evidence one can hope for. Toward that end, I appeal liberally to a variety of aspects of Kuhn's model, particularly those with direct connections to core features of evolutionary theory. Out of the many such connections that I exploit in the following chapters, perhaps the most important are the handful of social phenomena that Kuhn's model ties directly to science's ability to solve problems. And chief among these is the tendency of small groups of practitioners to effectively isolate themselves from other groups, as well as from nonpractitioners, as they slowly progress

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

toward a satisfactory understanding of nature. I examine this phenomenon in detail in chapter 3.

1.2.2 Karl Popper and Donald Campbell

Popper (1972) and Campbell (1974) both argue for a view about the development of scientific knowledge that centers on the role of the Darwinian mechanism of selective retention in the process of scientific theory change. The significant differences in their accounts are mostly differences of emphasis. Popper, perhaps unsurprisingly given his earlier work, focuses on the competitive dimension of the selection process:

> We choose the theory which best holds its own in competition with other theories; the one which, by natural selection, proves itself the fittest to survive. This will be the one which not only has hitherto stood up to the severest tests, but the one which is also testable in the most rigorous way. (Popper 1959, 91)

Campbell, for his part, stresses the idea that variation is random with respect to environment, or "blind": "I have gone beyond the others I claim to represent in my dogmatic insistence that all expansions of knowledge must involve a nonprescient variation and selective retention process at some level" (Campbell 1977, 505). Outside these different flavors, the philosophical model of knowledge production is essentially the same: practitioners propose theories, which they have designed to satisfy certain criteria. Some proposed theories do a better job of satisfying those criteria, and they are retained for future use. That future use will involve refinements to the surviving theories as new challenges arise. And so on, through a never-ending cycle of trial and error.

As with any philosophical theory, critics have raised a range of objections to each of these accounts. Negative responses to Popper's view focused on the fact that it is unnecessarily weird, invoking as it does a mystical "third world" where scientific knowledge resides (see, e.g., Bloor 1974), while those aimed at Campbell tend to complain about his "dogmatic insistence" on the blindness of conceptual variation (see Richards 1977 for an excellent analysis). I believe these criticisms are just in and of themselves, but I do not think they get at the issue of why the analogy in either Popper's or Campbell's hands fails to be very useful. Rather, something like the problem that concerns Godfrey-Smith is at work—namely, that we do not learn anything about the production of scientific knowledge that we did not know before viewing it through this evolutionary

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

lens. By itself, the idea that much of science proceeds on the basis of repeated trial and error does not exactly bring us to the threshold of profundity, and it's not clear what redescribing that process in Darwinian terms adds to our understanding of science. And although Campbell's conceptual and historical exploration of the mechanisms of trial and error is extraordinarily deep, that depth does not translate into a better understanding of the sense in which scientific knowledge is something that evolves in a Darwinian fashion, nor does his approach offer us any obvious avenues for learning something about science that we did not already know before we were encouraged to think of natural selection is just another kind of trial and error process.

The challenges faced by Popper's and Campbell's version of the evolutionary analogy are symptomatic of the highly restricted understanding of evolution that prevailed outside of evolutionary biology prior to the 1980s (or thereabouts). Not only is there an overly narrow focus on scientific analogies with *selection*, but the specific conception of natural selection in play is itself needlessly narrow. In Popper's case, the exclusive focus on competitive exclusion as the mechanism of selection occludes the interesting panoply of selective influences that can bear on members of a population, many of which have suggestive connections with similar pressures in the natural sciences. Likewise for Campbell's view: Apart from its being a prima facie implausible characterization of variation in scientific practice, blind variation is really just a limiting case in the theory of natural selection; natural selection can still operate in contexts where there is a nonrandom association between the demands of the environment and the properties of variants.

There are two important effects of employing such a tightly constrained conception of the evolutionary process. First, if we are motivated to make the analogy work, we are likely to be prone—as was Campbell—to misleading distortions in our representation of scientific practice (Godfrey-Smith's "shoehorning"). But more importantly, I think, we are discouraged from considering the potential for a structural match between science and evolution in all but a very few dimensions. The upshot is the suppression of an otherwise good analogy's power to generate new questions or suggest new avenues for research. There is so much more to evolution than the selective effects of competitive exclusion. And the fact that scientists select the theory that performs best is not something we needed the Darwinian theory of evolution to disclose.[2]

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

1.2.3 Stephen Toulmin

As an early philosophical venture into the social epistemology of science, Stephen Toulmin's *Human Understanding* (1972) deserves to be far better known than it is. Toulmin's inventory of the social factors that influence the nature and direction of scientific research is unapologetically thorough. It was the first large-scale systematic attempt to understand how nature, social norms, and epistemological ideals interact to produce the phenomenon of scientific knowledge in the wake of Kuhn's *Structure*. It is also admirably philosophical in its effort to extract from this milieu a conception of rationality that is both faithful to the complex multidimensionality of scientific decision-making and yet recognizably normative (a project that Kuhn never seemed to be able to treat with the seriousness that philosophers thought it deserved—or, anyway, not in the manner that they wanted it treated).

This is not the only sense in which Toulmin out-Kuhns Kuhn. Indeed, one gets the distinct impression that Toulmin grasped the epistemological and evolutionary implications of *Structure* far more quickly and firmly than Kuhn himself did—at least at the time he (Kuhn) wrote it; perhaps long after that. In particular, what Kuhn would only appreciate gradually and in chunks over the ensuing decades, and what Toulmin seems to have seen immediately, was that core social features of scientific practice tend to isolate and exert pressure on scientific communities such that they come to resemble biological populations in evolutionarily important respects. Although, as I describe throughout this book, Kuhn was clearly aware of the epistemic significance of isolation and "environmental" pressure, it was Toulmin who recognized the connection between the centrality of these phenomena on the one hand and the major conceptual features of "population thinking" on the other. Drawing that connection enabled Toulmin to see scientific disciplines as sensitive to a diverse range of influences—labeled an "intellectual ecology"—which together explained the content and transformation of scientific concepts over time.

The principal virtue of Toulmin's study is the way in which the application of population thinking facilitates a focus on the scientific community as the site whereupon scientific knowledge is modified by the evolutionary process; this is what compels him to embrace a view of how scientific knowledge develops that is uncharacteristically catholic for a philosopher of science at that time. To that extent, it is also for its time as thoroughgoing a social epistemology of science as one is likely to find anywhere. But as

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

an evolutionary framework for understanding science, it is only marginally less restrictive than those of Popper and Campbell. In the final analysis, there are three evolutionary concepts that matter for Toulmin's account: the scientific "population," the conceptual "variation" it produces, and the intellectual "environment" to which it is responding and by which it is ultimately shaped. There is no denying that these concepts are central to the theory of evolution, and the structural match between science and the evolutionary process they highlight is a significant advance over the approach taken by Popper and Campbell. By themselves, however, they are not able to give us more than a coarse-grained evolutionary image of the historical development of scientific knowledge. Also, they are not really invoked all that frequently. Evolutionary concepts are, for Toulmin, mostly used as suggestive themes for setting up a wide exploration of the social milieu of science; they are not primarily used to explain anything about science. Furthermore, as Godfrey-Smith might have observed, Toulmin's evolutionary framework again fails to disclose anything qualitatively distinctive about the production of scientific knowledge. As a causal framework for understanding science, essentially we are told that scientists have varying opinions, and that they make decisions on the basis of reasons that they believe matter to the growth of knowledge.

One last feature of Toulmin's study is instructive and worth noting. By the time Toulmin was writing, Kuhn had been given ample opportunity to articulate some sense in which the canons of scientific reasoning he described in *Structure* corresponded to a conception of epistemic rationality that could accommodate ideas of better or worse scientific reasoning at a more general remove than the highly localized and isolated contexts on which he had focused. In my opinion, Kuhn never really made a serious effort to combat the charges that his account was dangerously unepistemic in a variety of ways. I'd like to know why. Whatever the reason, Toulmin again seems to have perceived much more clearly than Kuhn did the importance of the problem and the need to treat it systematically. Toward that end, he (Toulmin) goes to considerable lengths to answer the philosophical demand to give an account of how the norms governing acceptable inferences can be fixed by local context while doing (some) justice to traditional notions of epistemic rationality. Although his bullet-biting answer probably would not have satisfied Kuhn's critics, his effort to treat the dominance of local epistemic norms in science as a constraint on a general conception of

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

rationality was a major achievement. (As a work of philosophy of science, I find *Human Understanding* to have been generally way ahead of its time.) But there is no real explanatory role for evolutionary considerations in Toulmin's account of rationality, and perhaps here again we see Godfrey-Smith's criticism of evolutionary analogies borne out: when the opportunity arises to make a distinctive contribution to our philosophical understanding of science, evolutionary theory is conspicuously absent.

1.2.4 David Hull

Like Toulmin, the connection between the evolutionary process and the way in which groups of researchers drive the development of scientific knowledge functions as the center of explanatory gravity in David Hull's *Science as a Process* (1988). Indeed, Hull himself seemed to suggest (12) that he was making good on Kuhn's wish that *Structure* had made more out of the evolutionary flavor of scientific communities. The specific focus in Hull's account is on *conceptual change*, which occurs through a selective process fueled primarily by competition among and cooperation within what he (I think appropriately) calls "demes." Internally, these demes foster the development of new ways of approaching the understanding of nature. Demes then compete with one another by critically testing each other's ideas. In Popperian fashion, theories that fail tests are selectively eliminated (378). Hull's analysis of selection, in that book and elsewhere, is structured around a relationship between two general kinds of entities: *replicators*, which make copies of themselves, and *interactors*, which interface with the external environment on behalf of replicators and affect their ability to make copies of themselves. In his account of conceptual change in science, the replicators are "elements of the substantive content of science—beliefs about the goals of science, proper ways to go about realizing these goals, problems and their possible solutions, modes of representation, accumulated data, and so on." These "elements of the substantive content of science" are transmitted and brought into confrontation with the conceptual environment by "books, journals, computers," and scientists, "the primary interactors in the conceptual development of science" (434).

Hull's project combines the strengths of both Toulmin and Popper-Campbell, while putting them on a richer evolutionary theoretic footing. Like Toulmin (and Kuhn), Hull sees the social structure of scientific communities as the primary driver of the development of scientific knowledge,

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

and both of them see these communities as performing an essential evolutionary role. Hull also views the process of natural selection as critical for understanding the growth of scientific knowledge. As in Campbell's framework, his account emphasizes the importance of inheritance mechanisms in science, since there is no evolutionary story to tell without them.

My authorial instinct at this juncture is to say, "Having said that," and then go on to summarize the numerous deficiencies in Hull's approach, which make my long-awaited intervention necessary. I'm going to resist that urge. Hull's book is superb. Nothing before or since has come close to it in the way it takes a very careful empirical look at the social dynamics of scientific practice and then connects them comprehensively to central questions in the philosophy of science. His deep reflection on the "credit" economy in science was justly influential, providing an essential and well-timed antidote to a developing strain of thought that threatened to reduce scientific debates to merely social disputes, rather than disputes tightly governed by epistemic norms.

Having said that, I should like to point to a few pathologies of Hull's account that hinder the generative potential of the evolutionary analogy. First, like his evolutionary predecessors, Hull's account is overwhelmingly focused on selection-driven conceptual change in response to the demands of an "environment" of social norms and competition that drive the scientific process. In this respect, his framework is nearly as restrictive as those that preceded him. While Hull was certainly the most sophisticated evolutionary thinker among those who pursued the evolutionary analogy, this sophistication is, regrettably, rather narrowly applied. And once again, the range of epistemic phenomena treated by Hull is nowhere near as broad as it could have been.

Another reason that the analogy suffers in Hull's hands has to do with his focus on the development of systematics as a case study. This is a highly nonrepresentative case. First, the competition between cladistics and other approaches to taxonomy was a pitched battle in which scientific meetings routinely devolved into shouting matches and even violence. Second, much of the fracas centered around philosophical principles about the nature of science whose adherents displayed a depth of commitment reminiscent of seventeenth-century Cartesian mechanists. While scientific debate can be quite intense, the excesses of the systematics wars are in no way characteristic of the social dynamics of science. This unfortunate period in the history

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

of science cannot, I fear, tell us much about the extent to which science in general resembles the evolutionary process.

A third major concern relates to the conceptual lens through which Hull views the evolutionary process. Hull, as I said, favored the replicator framework when it came to thinking about selection, and his preference is enthusiastically enforced in his analysis of science. Although I appreciate the power of this approach as a metaphysical framework for evolution by natural selection, I make no direct reference to replicators or interactors in the following pages, for a few reasons. The main reason is that I don't find it to be all that helpful as a way of fleshing out the evolutionary dimensions of the scientific process. I understand the appeal of its generality: rather than talk about analogies with the evolutionary process in biology, we can just talk at a single level about selection in general. But by doing this, we obscure potential connections with the evolutionary process. For one thing, very little theoretical work on evolution is framed in terms of replicators per se. If we are making a sincere effort to build out the analogy between evolution and science, it just doesn't make sense to frame the evolutionary process in terms of replicators and interactors when the entire edifice of evolutionary theory is built on a completely different conceptual foundation. The analogy's credibility depends to a significant extent on whether people familiar with the theory of evolution can easily recognize the version of the theory represented in the analogy. The replicator framework fails that test.

Generally speaking, the use of the replicator framework weakens our ability to use the enormous power, clarity, stature, and scope of classical evolutionary theory. It really only excels at representing the process of selection; and there is an enormous range of evolutionary phenomena that simply cannot be expressed in the language of replicators. And, as I have stressed throughout this discussion, an evolutionary analogy that restricts itself to the selective process is a deeply and needlessly impoverished tool for analyzing science. I hasten to emphasize, however, that this is not so much a problem with the replicator framework as it is a problem with the project of using evolutionary theory to illuminate the nature of scientific inquiry.

Let us conclude the Pummeling of the Predecessors by reflecting on what we've learned. This ritual beating has been necessitated by the fact that there needs to be a good reason why we should invest in further pursuit of the analogy between science and evolution. One general theme that has

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

emerged from the pummeling is that previous efforts in this domain have been unduly restrictive in their application of evolutionary theory. With this in mind, it is easy to see the justice in Godfrey-Smith's criticism. A narrowly conceived image of the evolutionary process provides very little in the way of opportunities to look for evolutionary themes in the practice of science and in the pattern of historical development of scientific knowledge. It is no wonder that the use of the evolutionary analogy has tended not to lead anywhere; it has not yet been given the chance to suggest pathways for inquiry that take us beyond the small set of epistemic phenomena that the analogy has traditionally targeted.

Another general theme disclosed by our survey suggests that the lackluster legacy of evolutionary models of science is due in part to their failure to highlight or solve any philosophical problems. To show that the use of the evolutionary analogy is philosophically important, we need to be able to see what sort of work it can do for us. Does thinking of science as an evolutionary system help to explain things about science that we didn't previously understand? Does it bring into view distinctive features of science that cry out for explanation? Does it shed light on philosophical problems beyond science, such as those associated with epistemology more generally? No demands of research are satisfied by an analogy that offers nothing except for a redescription of well-known phenomena.

I'm going to use these two general themes to dictate the structure and direction of the book's argument. Broadly speaking, my strategy involves three components. First, I try appeal to as broad of a range of evolutionary ideas as I can (with due consideration for readers' patience). For reasons to be explained shortly, I maintain a prominent place for natural selection. However, in the framework I develop, selection is something that appears in greater or lesser degrees, depending on other conditions which themselves are part of the evolutionary process and form part of the theory of evolution. As I will show, the relative *in*significance of selection in certain contexts is crucial to the explanation of phenomena that we know are important to the growth of scientific knowledge. Second, I subject to evolutionary analysis as wide a range of scientific practices and historical patterns as space and systematicity will allow. I have the good fortune of writing at a time in which many trends associated with the practice of science have been disclosed by previous research, primarily in the history and sociology of science. Insofar as these trends are genuine, I take them to be

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

part of the growth process for scientific knowledge. As such, they function as data with which an epistemology of science must reckon. In Kuhn's memorable words: "How could history of science fail to serve as a source of phenomena to which theories of knowledge may legitimately be asked to apply?" (1962b, 9).

Third and last, I attempt to motivate the perspective that a number of features of scientific inquiry only really make sense in the light of evolutionary theory. Chief among these is the apparent tendency of modern scientific inquiry to make progress. Although this tendency is more or less taken for granted by most scientists, most casual observers, and many philosophers, it is in no way a fact generally acknowledged among historians, sociologists, and anthropologists of science. As I explain throughout the book, I do think that science makes progress in an important and recognizably epistemic sense. But I am sympathetic to much of the skepticism surrounding the notion of scientific progress. Simply put, why would it be that human beings make progress when they practice science but scarcely otherwise? An answer that has proved immensely attractive in some disciplines is that the appearance of scientific progress is illusory; human scientific endeavors actually embody the same lack of progress as do our nonscientific pursuits. Viewed through this lens, "no progress" should be the default expectation for scientific inquiry. The fact that it appears to routinely deviate freakishly from this default expectation is a puzzle.

Part of solving this puzzle will involve indulging in the same fascination with the resemblance between adaptation and scientific problem-solving that captured the attention of previous investigations of the connection between science and the evolutionary process. But, as with any indulgence, the rush to treat science as an intrinsically adaptive system must be tempered by the discipline to accept that science just as routinely does not behave in the way that we would expect from a system governed purely by natural selection. Now, because these nonadaptive features also form part of the process by which scientific knowledge grows over time, they too must thereby "serve as a source of phenomena to which theories of knowledge may legitimately be asked to apply." That is, any account of scientific knowledge that leans heavily on evolutionary adaptation must also be able to accommodate aspects of the production of scientific knowledge that clearly do not correspond to the adaptive process. Solving the puzzle of progress by appealing to adaptation generates a puzzle of its own—namely,

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

how to understand aspects of epistemic growth that have traditionally been viewed as progress but not for reasons related to the adaptive process.

1.3 From Progress to Directionality

Stephen Jay Gould recognized that "our obsession with progress records something larger, deeper and vitally important," "a crucial generality that we must pursue—the study of directional change in history" (1988, 319). For complicated reasons, Gould and others are reluctant to describe directional change as progress. But directional patterns are central to the study of evolution, and Gould rightly worried about conflating hostility toward progress with a blanket denial of evolutionary directionality. He thus helpfully recommended replacing the notion of progress in evolutionary contexts with the "operational notion of directionality." Similarly, whether or not one is disposed to perceive in science a tendency to move inexorably toward epistemically superior states, it is undeniable that scientific inquiry appears to exhibit a strong directionality at multiple levels; even the nonepistemic sociocultural models developed by contemporary historians and sociologists of science do not dispute this basic point.

At the most general level, and in keeping with traditional notions of scientific and evolutionary progress, we are interested in the historical pattern whereby science appears to move inexorably toward increasing degrees and kinds of specialization. More locally, we will be concerned with the apparent tendency within science to move—often equally inexorably—toward particular ways of understanding and investigating nature. Both sorts of patterns—the general increase in depth and breadth on the one hand and the highly constrained pattern of refinement within a given research tradition on the other—are closely associated with traditional notions of scientific progress. Both sorts of patterns possess the kind of directionality that has replaced notions of progress in biological contexts.

Since models of scientific change are concerned fundamentally with the empirical reality of science's apparent directional tendency, and only secondarily with articulating the sense in which that tendency is a characteristic of epistemic progress, I propose that we first narrow the focus of our investigation by restricting the explanatory purview of these models to these historical patterns. I propose to divide the investigation into two distinct types of directionality: (1) a *vertical* type, characterized by increasing

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

degrees of refinement or specialization, and (2) a *horizontal* type, characterized by increasing *kinds* or *branches* of specialization.

A successful explanation for the apparent historical tendency toward refinement and branching within scientific inquiry will reveal the nature of what it is we refer to as "scientific progress." And there is no guarantee that, once we understand what causes this apparent pattern, the referent of "scientific progress" revealed by our investigation will correspond in any way to our image of an activity that predictably and naturally moves to better and better epistemic states. Where we find that science does not reliably move to epistemically superior states, it would be prudent to deny that there is such a thing as scientific progress in those cases. Only when special conditions are realized will the apparent directionality of science amount to the progress of science.

The remainder of this chapter involves an effort to spell out in detail what those special conditions are. One part of that effort requires being able to distinguish between real and merely apparent directionality. But real directionality will not be sufficient to secure the conclusion that science is progressive in a straightforwardly epistemic sense. We will have to show in addition that the sort of directionality science exhibits is of an epistemic sort. This, too, has been an item of significant controversy in the years following Kuhn's *Structure*.

1.3.1 Pseudodirectionality

Let us look at a couple of different ways in which the apparent directionality of science might fail to support the idea that science makes progress. One way would be if the directionality of science is *merely apparent*; what appears to be directionality dissolves when subjected to serious scrutiny. Call this *pseudodirectionality*. A classic description of pseudodirectionality is provided by Kuhn in chapter 11 of *Structure*, "The Invisibility of Revolutions." There Kuhn argues that, for "functional reasons," the history of science is presented in science textbooks, in science classes, and in popular accounts as one of cumulative linearity: Each time t in the history of a line of inquiry is depicted as the direct outcome of whichever research at $t-1$ most closely resembles what we eventually came to believe at t, which is itself historically relevant only because it most closely resembles what we currently believe. The appearance of directionality is generated in two steps: first, by emphasizing only those aspects of earlier efforts that seem

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

to have been preserved in some form; and second, by ignoring or suppressing any research in the history of science that cannot be credibly made to look like an underdeveloped antecedent to current scientific consensus. Thus, celestial mechanics is made to look as though it proceeded directly from Kepler to Newton by (1) focusing on the logical equivalence between Kepler's laws and the laws that inspired Newton's law of gravitation and (2) failing to mention the aspects of Kepler's work that Newton rejected or ignored, such as his mystical mechanical explanation for how the planets are propelled around the Sun (the *anima motrix*). Since this is all the history of science that a student needs to know in order to understand current scientific consensus, there is no need to waste time with messy, scientifically irrelevant details. Out of the historical morass of numerous researchers blindly groping for an answer, we manufacture a tidy linear trek from one lone genius to the next, thus creating the impression of directionality by a selective manipulation of the line of inquiry's actual causal history.

If the development of science is only pseudodirectional, then there is no historical evidence for a causal process that produces increases in epistemic well-foundedness through refinement of ancestral states of scientific inquiry. Rather, what is actually happening in pseudodirectional cases is that a scientific community first changes its mind and then looks for a direct path between the now-established consensus and some previous view. In Kuhn's memorable words:

> Partly by selection and partly by distortion, the scientists of earlier ages are implicitly represented as having worked upon the same set of fixed problems and in accordance with the same set of fixed canons that the most recent revolution in scientific theory and method has made seem scientific. No wonder that textbooks and the historical tradition they imply have to be rewritten after each scientific revolution. And no wonder that, as they are rewritten, science once again comes to seem largely cumulative. (Kuhn 1962b, 138)

As depicted in figure 1.1, in the event, the community's change of mind may actually cause absolute epistemic well-foundedness to decrease. The appearance of increase is then generated by locating pseudopredecessors within ideas held in the scientific past that happen to have not been as epistemically well founded as the current consensus views. On this understanding of apparent directionality, progress is a myth cultivated by conveniently arranging and rearranging historical facts whenever the scientific community changes its mind and then appealing to a fictitious causal process of

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

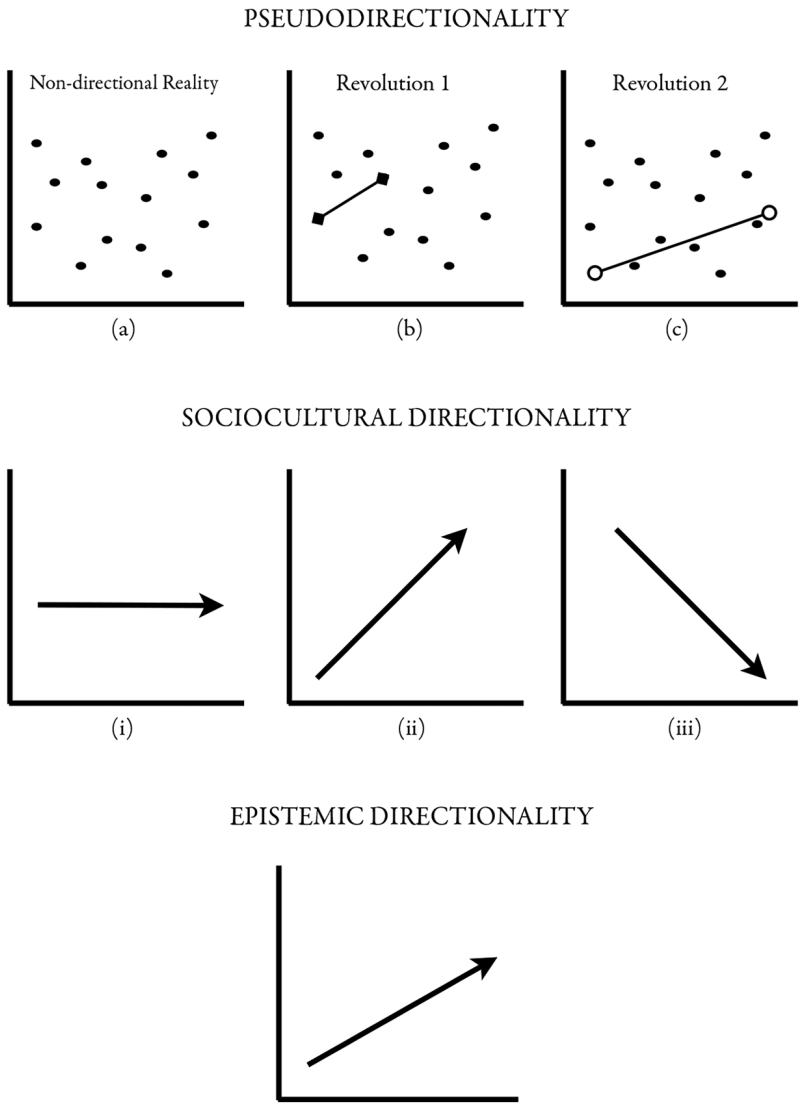

Figure 1.1
The directionality of science. Each graph plots epistemic well-foundedness as a function of time under different types of pressures. *Pseudodirectionality*: (a) the history of science is actually non-directional; (b) after Revolution 1, the historical development of ideas is depicted in a way that depicts current views as directly related to some previous views; (c) after Revolution 2, newly adopted views are depicted as direct descendants of *different previous views*; as in (b) and (c), overall epistemic well-foundedness might decrease from one revolution to another. *Sociocultural Directionality*: scientific knowledge is responding to pressures that might vary orthogonally to epistemic well-foundedness, which allows it to exhibit a range of relationships to epistemic well-foundedness: (i) flat; (ii) increase; (iii) decrease. *Epistemic Directionality*: scientific knowledge, responding to epistemic pressure, exhibits a consistent rise in epistemic well-foundedness.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

refinement to explain the manufactured historical sequence. In fact (claims this account), the only causally influential process in the historical development of scientific ideas is the scientific community's discontinuous and somewhat sporadic changes of mind. Since those changes are not actually governed by any causal process that reliably generates growth, there is no basis for thinking that science itself is inherently progressive.

1.3.2 Sociocultural Directionality

Paul Forman's (1971) study of early quantum theory is a classic example of a model of scientific inquiry that depicts the historical development of a science as at once directional and nonepistemic. Forman carefully builds an argument supporting the conclusion that the profound transformation of physical theory that took place during the early twentieth century was caused by an antecedent shift in more general cultural attitudes that occurred in Germany in the aftermath of World War I. He argued that "extrinsic influences led physicists to ardently hope for, actively search for, and willingly embrace an acausal quantum mechanics" (1971, 3)—specifically, that the transition from a causal-mechanical approach to physical processes to one in which causal notions played little or no role was fueled predominately by German physicists' and mathematicians' attempts to "adapt" to the prevailing intellectual hostility toward causality and causal notions generally.

The epistemologically significant point here is the extent to which the direction of theoretical physics in Germany seems to have been so sensitive to and, consequently, so strongly influenced by broader cultural concerns that are otherwise far removed from the narrow agenda of a small community of specialists trying to generate scientific knowledge. In this way, Forman's study suggests that the particular path along which scientific inquiry develops may be largely in response to "the intellectual currents of the day"—features of social, historical, or cultural context that are not directly connected with the need to refine scientific theories, questions, and techniques in order to arrive at an improved understanding of nature; indeed, in this case, "intellectual currents which . . . were fundamentally antithetical to the scientific enterprise" (1971, 109). Here, then, is an instance where the most advanced scientific thought of the day was allegedly undergoing a process of directional refinement structurally akin to progress, but in response to considerations that were not specifically tied to developing a better understanding of natural phenomena.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Forman makes no attempt to deny that the quantum mechanics that developed out of that period was and has continued to be enormously successful. But one can see how the scientific community's susceptibility to developing inquiry in a direction more consonant with the prevailing cultural winds might render extremely fragile any natural potential within science to move toward improving scientific knowledge. For, were it the case that science is characteristically sensitive in this way, the actual direction of scientific inquiry might then easily move orthogonally to any of the trajectories that increase either the degree or the number of kinds of specialization of scientific knowledge. Just as we would expect a morphological trait evolving in response to pressure from environmental factor *X* to vary randomly with respect to factor *Y*, the epistemic null expectation for directionality generated by sociocultural pressure is that the development of science will vary randomly with respect to improvements to our knowledge; perhaps our theories will become less accurate, or perhaps the breadth of our knowledge will contract rather than expand. And if sociocultural directionality gives us no reason to expect scientific inquiry to reliably move toward increases in depth and breadth, inquiry may then develop directionally by carefully refining scientific theory and practice, and yet do so without actually making progress.

1.3.3 Epistemic Directionality: The Essence of Scientific Progress

I have described adaptation to sociocultural pressures as simultaneously (1) adaptation to pressures that are external to science and (2) adaptation to pressures that are nonepistemic. This is no accident. As I argued above, part of the reason sociocultural adaptation qualifies as nonepistemic is because it might easily vary randomly with respect to increases either in degree or kinds of specialization. To put it another way, if and when inquiry develops in response to sociocultural pressures, then the direction in which inquiry develops is in general not explained by improvements to our knowledge. Rather, it is explained by adaptation to pressures that vary randomly with respect to epistemic goals.[3]

This argument suggests the following principle regarding what qualifies as an epistemic pressure: *A pressure is epistemic if and only if adaptation to it is explained by increases in the depth or breadth of our knowledge*. Call any increase in depth or breadth an *improvement*. The motivation for characterizing an epistemic pressure in this way is straightforward. It follows from

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

the more general notion of ecological pressures like, for instance, predation. To say that some ecological factor is a predation pressure is to say that it causes populations to evolve traits that improve their ability to avoid predation. Some of these traits may, by coincidence, later prove to produce other kinds of benefits. But the spread of these traits throughout the population is not explained by those other benefits. Likewise, we say that some factor is an epistemic pressure when it causes us to develop improvements in our knowledge. Some improvements may, by coincidence, later prove to produce other kinds of benefits. But the development of those improvements to our knowledge is not explained by those other benefits.[4]

The argument above also implicitly asserted that sociocultural pressures are not "internal" to science. Is there a plausible basis for this assertion? Let us define a pressure as being *internal to science if and only if adaptation to it is explained by achievements that are recognizably important to most practitioners, but not to most nonpractitioners*. I adopt the commonsense assumption that there are some things that are cared about by most practitioners in a field but not cared about by most nonpractitioners. Sometimes practitioners won't care, and sometimes nonpractitioners will, but the relative size of these groups tends to be small. I also adopt the commonsense assumption that, for the relevant set of achievements, the shared sense of importance held by most practitioners is explained by the fact that they are practitioners, and the shared sense of unimportance or bemusement held by most nonpractitioners is explained by the fact that they are not practitioners. Call such pressures as satisfy these conditions *internal*. An example of an achievement that responds to an internal pressure would be something like the article published in *Physical Review* on the conditions under which a sequence of events is reversible (Bonança and Jarzynski 2016). This result will probably be regarded as important by most statistical physicists and other practitioners of closely related disciplines, but it will fail to attract the interest of more than a negligible fraction of nonpractitioners.

Thus, if, as I have assumed, the pressures that are internal to science are epistemic, then it ought to be the case (combining the "internal" and "epistemic" definitions) that (1) improvements to the depth or breadth of scientific knowledge are the kinds of achievements that are important to most practitioners, but not to most nonpractitioners, and (2) adaptation to pressures internal to science is characteristically achievable only through improvements in scientific knowledge. These results correspond closely to

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Kuhn's observation that "the unit of scientific achievement is the solved problem," which is significant, given their independent derivation (Kuhn 1962b, 169). Only under these conditions can scientific knowledge be said to exhibit a directional tendency toward improvements in scientific knowledge.

We began with the core idea expressed by a variety of thinkers that the evolutionary nature of science is what makes science epistemically special. Through successive refinements, we arrived at a considerably more precise, more analytically approachable thesis, which we can break down into the following components:

To explain the epistemic progress of science using the principles of evolutionary theory is to show that scientific progress is the expected result of the evolutionary nature of science.

Like evolutionary progress, scientific progress can be essentially understood in terms of its *directionality*.

Directional advances in scientific knowledge amount to increases in specialization. These increases can be of two kinds: (1) *Vertical*: increases in degrees of specialization or refinement; (2) *Horizontal*: increases in the range of phenomena under investigation.

In order for the directionality of science to qualify as progress in an epistemic sense, that directionality cannot be an artifact of the narrativization of the history of science (*pseudodirectionality*).

In order to qualify as progress in an epistemic sense, the directionality of science needs to be an expected result of pressure specifically to improve scientific knowledge.

Together, these components effectively express what it means for scientific progress to be explained by its evolutionary nature; at any rate, that is what I will take it to mean for the purposes of this project. The hard work ahead involves supporting the general idea that scientific knowledge actually exhibits the signatures of an entity undergoing directional selection in the sense described. Before turning to that hard work, we can add a few more constraints that will allow us to better structure our investigation. The relevant constraints are general ones governing what sorts of entities can possibly undergo evolutionary adaptation. What sorts of general properties do scientific populations need to possess in order to even be candidates for directional selection?

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

1.4 Scientific Practice as a Darwinian Population

There are, then, two straightforward senses in which the evolutionary process exhibits progress. First, populations progress toward local phenotypic optima. Maybe they get there, maybe not. But the fact is that selection pushes them in a particular direction in the space of physically possible morphologies ("morphospace") toward a point at which fitness is maximized in a given environment. This is the process of adaptation. Second, the inventory of taxa in a region tends to increase to a point of saturation or "carrying capacity." Again, whether it gets there or not, there is a tendency built into the evolutionary process toward increasing the number of taxa so long as there is energy to be consumed. This tendency, which Darwin in the *Origin* labeled the "Principle of Divergence," is—in a sense to be articulated in chapter 5—a by-product of the selective process.

As a first approximation, these two senses of evolutionary progress correspond to the tendency for scientific knowledge to display consistent growth in both its depth and its breadth with the passage of time. Now if, as I have just suggested, the increasing depth and breadth of scientific knowledge are both consequences of the selective process, then scientific knowledge has to be the sort of thing that can be subject to selection. Not all group structures are equally capable of supporting the sort of directional trends apparent in science, and still fewer are capable of undergoing the distinctive subvariety of directionality known as adaptation. An explanation for the directionality of science that appeals to the evolutionary process, and more specifically to adaptation, needs to provide independent evidence for scientific knowledge's satisfaction of the general constraints that allow populations to evolve adaptively.

To marshal this evidence, I appeal to Godfrey-Smith's highly structured yet flexible notion of a *Darwinian population* to describe the kind of entity that can evolve by natural selection, and thus, the kind of entity that scientific knowledge will need to be if it evolves adaptively. These populations are characterized in the following way:

> A *Darwinian population in the minimal sense* is a collection of casually connected individual things in which there is variation in character, which leads to differences in reproductive output (differences in how much or how quickly individuals reproduce), and which is inherited to some extent. (Godfrey-Smith 2009, 39)

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Different Darwinian populations exhibit different degrees of susceptibility to evolution by natural selection, depending on the extent to which they possess each of three central components:

Heredity (*H*): the extent to which offspring resemble parents, due to the causal role of parents;

Intrinsicality (*S*): "the extent to which differences in reproductive output in a population depend on intrinsic features of the members of the population";[5]

Continuity (*C*): the extent to which the magnitude of trait differences tracks the magnitude of their fitness differences.

In Godfrey-Smith's account, "paradigm" Darwinian populations—those ripe for evolution by natural selection—are those that exhibit high-*H*, high-*S*, and high-*C*.

In addition to outlining general conditions on what makes a population *Darwinian*, Godfrey-Smith described how those conditions might be realized when the population was made of things like "habits, accents, and ideas"—what I will refer to generally as *practices* (2009, 153). Practices reproduce by causing more instances of themselves, such as when my neighbors adopt a cross-cut style of lawn mowing based on how I mow my lawn; my lawn-mowing style is the "parent" of theirs. These several instances of the cross-cut lawn-mowing style form a lineage, since each instance can be traced back to a "parent" instance that is causally responsible for (1) its existence and (2) its resemblance to its parent(s) (158).

A lineage of practices constitutes a Darwinian population whenever differences in practice-instances cause differences in the reproductive propensities of those practice-instances—that is, in their propensities to cause more instances of themselves. One of my neighbors decided to modify my original cross-cut style by cutting diagonally rather than side to side. Because of its greater aesthetic appeal, more neighbors tend to adopt the modified diagonal cross-cut style than they do the side-to-side style. Consequently, the majority of lawns in our neighborhood are now diagonal cross-cut. Although each lawn in the neighborhood is now cut in a style descended from my initial cross-cut, the population of styles has been modified by natural selection for aesthetic appeal. Through this selection process, the Darwinian population consisting of neighborhood lawn-mowing practices has evolved from side-to-side cross-cut to diagonal cross-cut style.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

I have chosen the general term "practices" to refer to the variety of things implicated in the constitution and growth of scientific knowledge. Under this rubric fall such apparently varied notions as experimental techniques, concepts, metaphysical commitments about the fundamental constituents of reality, canons of inference, substantive scientific theories, and research problems—in short, the panoply of factors that distinguish the practice of science and the associated production of scientific knowledge. My characterization of this panoply descends from Kitcher (1993), who analyzed scientific progress in terms of "practices" that are passed down with modification from "veterans" to "apprentices." I adopt the term "practices" because it (1) is very general, (2) has no insurmountably unintuitive connotations, and (3) has been used by previous analysts to do the sort of conceptual work for which I intend it. I use the singular term "scientific practice" to refer to the collection of such practices.

To this extent, the model I present in the book can be described as *practice-oriented* and can be viewed as a beneficiary of the "practice turn" in the philosophy of science that has become increasingly mature over the past few decades. Although the models presented by Toulmin, Hull, and Kitcher make important use of aspects of practice outside of theory, their focus is still strongly centered on the modification of scientific concepts and theories as the key to understanding the historical development of scientific knowledge. More recent work by such practice-oriented philosophers of science as Hasok Chang, Emily Grosholz, Alan Love, Hans-Joerg Rheinberger, Joseph Rouse, Ken Waters, and Andrea Woody have endeavored to show that, in fact, theory is only one of many drivers of scientific change. Perhaps unsurprisingly, these philosophers draw substantively on the history of science in support of their efforts to widen the scope of philosophical analysis.[6]

Some might object to, for example, characterizing scientific theories or research problems as "practices." In general, there is no agreement on how best to characterize this range of phenomena, and there probably never will be. But nothing said below hangs substantively on the terminology here; one might just as well use the generic term "φ" to capture the relevant sort of phenomenon. The important points are that (1) each φ is employed by practitioners at some level in the course of inquiry and for its sake; (2) the set of φs includes, perhaps primarily, the aspects of scientific inquiry that cannot credibly be categorized as "theory"; and (3) it is the evolutionary

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

modification of these φs that produces the directional pattern we associate with scientific progress.

Another concern might be whether we can individuate practices in a way that allows us to see clearly where one practice ends and another begins or to distinguish between an instance where a practice is modified, as opposed to an instance where it is abandoned and replaced.[7] For example, is the change from cross-cut to *diagonal* cross-cut an instance of modification or an instance of replacement? I view this problem as analogous to the problem of behavioral character individuation in evolutionary biology. We observe that a characteristic mating dance is widespread across males in a certain species of bird. We further observe that a handful of males, while exhibiting a broadly similar dance, insert a "bow" where most males do a "twist." Are these two different mating dances, or is the bow a mere modification of the canonical dance? Although the problem is a serious one, it has not prevented us from learning an awful lot about how populations of organisms evolve (Mesoudi et al. 2013, 210). For example, we do not need to answer this question in order to inquire into whether the relative frequencies of bows and twists are changing from one generation to the next. Analogous concerns are part of the very fabric of paleontology. As with most issues in classification, which of these options ends up being the best answer to this question will often turn out to be which answer allows us to satisfactorily organize our knowledge of the focal system and explain its behavior over time. I think the least productive route we could take would be to let these kinds of worries impede a philosophical investigation, as they quite appropriately have not held up the investigation of evolutionary biological change. New and better understandings of how to individuate characters are always welcome but the show must go on.

Applying the general characterization of a "paradigm" Darwinian population to the specific case of an evolving population of practices entails that a population of practices is a "paradigmatic" Darwinian population whenever:

Practices reproduce with high fidelity; "parent" instances will strongly resemble "offspring" instances.

Differences in practices' reproductive success depend on intrinsic features of the practices themselves; some practices are just intrinsically better suited than others to a given environment.

The magnitude of differences between practices tracks the magnitude of their fitness differences; small differences between practices should, ceteris paribus, translate into small fitness differences between practices.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Conceived of as a collection of practices, the susceptibility of scientific knowledge to evolution by natural selection depends on the degree to which scientific practice instantiates these three ideals.

As a tool for drawing an analogy between science and evolution, much of the power of the "Darwinian populations" framework lies in its ability to represent a natural selection's *degree of influence* over a population as a function of that population's degrees of heredity, intrinsicality, and continuity. Different scientific disciplines possess these properties in different degrees, as do particular disciplines at different times. As I will argue, those differences are reflected in the differing degrees to which scientific practice in those disciplines is refined by natural selection. In turn, these differences track differences in the propensities of disciplines to make progress in the manner specified above.

The ability to represent natural selection's degree of influence over a population's evolutionary trajectory includes the ability to represent instances in which that influence is negligible. In the chapters that follow, I will argue in many places that there are aspects of the growth of scientific knowledge which are best understood as instances of this kind. In these cases, the use of the "Darwinian populations" framework is particularly valuable because it is a natural way of depicting phenomena that are (1) epistemically significant and (2) evolutionary but (3) not easily understood in terms of increasing adaptedness. Because the framework allows us to set the heredity, intrinsicality, and continuity "dials" to different points on a continuum, it makes it easy to see how scientific practice is modified by the evolutionary process even when the modifying role of natural selection is relatively insignificant. Evolution does not cease to operate simply because natural selection is nudged out of the driver's seat. Understanding the nonadaptive evolution of scientific practice under certain conditions is an important virtue of which previous uses of the evolutionary analogy did not have the resources to avail themselves.

1.5 Science as a *Cultural* Darwinian Population

The capacity to represent other styles of evolutionary change is likewise encouraged by the fact that, as Godfrey-Smith shows, the notion of a Darwinian population can be expanded to represent the evolution of human culture, conceived of as a population of practices (Godfrey-Smith 2009, chap. 6; Godfrey-Smith 2012). In the introduction I briefly alluded to the

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

idea that modern science is a culture characterized by a highly idiosyncratic set of social norms, norms that have turned out to be immensely powerful for promoting the growth of knowledge. This suggests another way in which the "Darwinian populations" framework is well suited for building out the analogy between evolution and the scientific process. Just as the concept of Darwinian population gives us a simple framework for depicting a rich array of evolutionary processes, thinking about the historical development of scientific knowledge as an instance of cultural evolution provides us with yet another source for mining evolutionary analogues to natural science. Where the notion of a Darwinian population will allow us to bring scientific knowledge into contact with the full range of evolutionary phenomena, the theoretical resources of cultural evolution give us ways of representing evolutionary features of scientific knowledge that reflect its status as a specific form of human cultural expression. There are aspects of the evolution of human culture that the cultural evolutionary framework treats as particularly salient (e.g., imitation) but are of marginal significance within the broader context of classical evolutionary theory. Looking at the development of scientific knowledge from the perspective of human cultural evolution thus brings to the fore opportunities to analogically connect science and the evolutionary process that would be difficult to raise naturally through a standard evolutionary approach.

Exemplary of this fact, and in keeping with the aim of bringing the full gamut of evolutionary processes into the fold, is the relatively significant emphasis on individual learning in the cultural evolutionary framework. Although individual learning features in a modest variety of different corners in evolutionary theory (e.g., in bird song, or mate choice), it is by no means a canonical feature of the evolutionary process. By contrast, individual learning plays a central and distinct role in cultural evolutionary theory, as it is naturally presumed to have a pronounced influence on the evolution of culture. This asymmetry is important on account of the fact that the prevalence of individual learning tends to weaken the influence of selection. As Mesoudi et al. (2013, 203) observe, "It seems likely that nonselective change due to individual learning . . . [is] more significant in cultural evolution than are nonselective processes in genetic evolution." Thus, the use of a cultural evolutionary framework can facilitate the explanatory power of nonselective mechanisms of evolutionary change that might have otherwise been occluded by a classical evolutionary approach.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Beyond the general expansion of the evolutionary toolkit, research on the evolution of specific cultural practices can and should be used as template for thinking about how to structure evolutionary investigations of particular scientific practices. Over the past couple of decades, Mark Pagel has developed a research program on the evolution of languages that amply illustrates how fruitful the use of evolutionary theory can be when its power is turned toward problems associated with human cultural change. In Pagel's framework, a word is a practice that is used for the purpose of reference. The spread and descent of a word is determined by a combination of social selective pressures on linguistic practice, demographic and geographic features of the word-using populations, and time. Through the use of phylogenetic and statistical tools, Pagel has shown how these determinants can be used to very successfully model the history of languages as a series of evolutionary trees of "discrete heritable units" of the sort that compose languages (e.g., words and syntax) (Pagel 2009, 406).

I view Pagel's research program as the current high-water mark for the sort of evolutionary approach to specific cultural practices envisioned by Toulmin, Hull, and Kuhn. No one who encounters this work could fail to come away with the impression that evolutionary methods are both appropriate and extremely powerful ways of understanding changes in human practices over time. Although the evolutionary histories of the concepts of interest to philosophers and historians of science are for many reasons more difficult to track than the histories of words themselves, the two phenomena are obviously related, and the notion of a concept as a "discrete heritable unit" forms much of the basis for philosophers' appeal to the evolutionary process. As such, it would seem appropriate to employ something like Pagel's phylogenetic approach to the historical development of scientific concepts—indeed, of scientific practices more generally.[8] There are many excellent conceptual histories on topics in science and mathematics that could, in principle, support such analyses. Regardless of whether a phylogenetic conceptual history could actually meet the standard set by Pagel, however, the spirit of his approach exemplifies what it means to take seriously the evolutionary nature of cultural change. To the extent that the development of scientific knowledge is such a change, an approach like Pagel's represents what is possible when we conceive of scientific knowledge as governed by the evolutionary process.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

2 The Persistence of Practice

Offspring can resemble their parents to varying degrees, and the degree of resemblance between generations strongly affects the population's ability to solve the problems in its environment. There is a fairly straightforward sense in which a high degree of resemblance between parents and offspring is required for them even to share the same selective environment. Because the selective environment—or a relevant portion of it—must remain stable in order for directional selection to operate, it follows that populations can only undergo directional selection in the case where offspring closely resemble their parents and when offspring environments closely resemble parent environments. This basic assumption of evolutionary theory is reflected, for example, in the conception of fitnesses as descriptive of a particular genotype in a particular environment (Brandon 1978).[1]

If a trait is well suited for a given environment, then that trait's ability to make high-fidelity copies of itself facilitates its persistence in that environment. If, however, a trait is not well suited for that environment, or if that environment is subject to change, making high-fidelity copies of itself won't help much in the long run. The adaptive value of high-fidelity copying in general depends on two factors: (1) a trait's distance from a phenotypic optimum and (2) the stability of the environment. If a trait is optimally designed for its environment, that trait's ability to copy itself perfectly is favored so long as the environment remains stable. But if a trait is suboptimal, perfect copiers do poorly relative to imperfect copiers because the latter can explore morphospace in pursuit of the environment's stable optimum. As the trait nears an optimum, exploring morphospace becomes increasingly risky, and so there is pressure to increase the fidelity of its copying. At the optimum, copying should be perfect. For similar

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

reasons, perfect copying is ill favored in changing environments. An optimized trait is optimal only with respect to a certain set of environmental conditions. Once those conditions change, the optimum begins to recede in morphospace, and pursuing the shifting optimum requires that copying be somewhat less than perfect.

The foregoing suggests a general principle governing systems that undergo directional refinement: when a trait is optimized, there will be pressure to both suppress variation (i.e., to perfect the trait's copying) and stabilize the environment. Of these two pressures, the pressure to stabilize the environment has to be primary.[2] Refinement only makes sense relative to a particular set of environmental demands. Whether high-fidelity copying facilitates refinement in turn depends on the environmental demands remaining relatively stable. If these demands are in constant flux, neither high- nor low-fidelity copying will help to promote the persistence of a trait. Perfectly copied traits ensure their faithful representation only if they are copied, but whether they are copied at all depends largely on how well suited they are for the environment; perfectly copied lethal traits do not persist.

These are some of the general selective mechanisms that causally promote the stability of a biological population's selective environment and the relative stability of phenotypes over time within that environment. Paradigmatically Darwinian populations—populations whose change is dictated purely by natural selection—copy their traits (including their environments) with near-perfect fidelity. Idealized epistemic communities follow suit: their practices—methods, concepts, research problems, and so on—are passed from veterans to apprentices with minimal variability. Differences in actual epistemic communities' susceptibility to modification through selection is (partly) a function of differences in how closely they approximate this ideal. In turn, these differences track differing degrees of directionality.

The dependence of directional refinement on practice copying is by no means a unique feature of the natural sciences. Craft traditions, from which the veteran-apprentice model ultimately derives, develop directionaly in precisely this way (Smith 2004; Long 2011). Kuhn makes a brief allusion to the connection between paradigm-governed scientific communities and the premodern "schools of thought," which were known for perpetually refining a corpus of ideas, typically through the memorization of key texts that formed the basis of subsequent reflection and insight (see, e.g., Makdisi

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

1981). These communities formed what Karin Knorr Cetina aptly dubbed *epistemic cultures*, and certain of their well-documented traditions can be mapped onto a Darwinian population framework in a way that explains why aspects of these cultures seem to display signature features of the selective process. In this chapter, I look at some of the cultural traditions we find in the sciences that tend to promote a strong resemblance in practice across successive generations of scientists. I describe and illustrate a multitude of mechanisms by which scientific practices come to be inherited, assessing their degree of fidelity and venturing an account as to why the inheritance mechanisms take the particular forms that they do. This analysis will help us understand the ways in which scientific practice might come to resemble the kind of population that is apt to exhibit the directionality we associate with populations evolving under natural selection.

The notion that scientific practice is passed down from veterans to apprentices is certainly not a novel idea, nor is the observation that high-fidelity copying of practices is actively pursued by members of some scientific communities. Furthermore, as I relate in what follows, Kuhn invested this cultural norm with the gravest epistemic significance. He understood it not only as the defining feature of scientific communities but also as a necessary causal ingredient in the process of scientific discovery. What he seems not to have grasped when he wrote *Structure* was that this tradition of copying is also significant from the perspective of evolutionary theory; it was left to later commentators interested in the analogy between science and evolution to draw this connection. But although it is well known that earlier generations of scientists influence later ones, I do not think the evolutionary explanation for how this tradition causally underlies a certain form of scientific progress has been sufficiently appreciated. The use of the Darwinian populations framework makes it clear why this causal relationship should exist. The contribution I hope to make here is to use the framework of Darwinian populations to further refine our understanding of the epistemic significance of high-fidelity practice copying. This framework gives us the ability to clearly articulate the way in which high-fidelity copying (partly) determines the susceptibility of scientific practice to the form of scientific progress associated with directional refinement—the *vertical* form corresponding to the deepening of scientific knowledge.

In that connection, the Darwinian populations framework also allows us to make sense of why actual scientific communities show such marked

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

discrepancies in their propensity to undergo vertical progress. Where such discrepancies exist, their source might ultimately lie in the relative strength with which respective communities insist on high-fidelity copying. Toward this end, the chapter concludes with a detailed comparison of the abstract model and the well-studied episode of chemist Justus Liebig's development of the modern research group in the mid-nineteenth century. The defining feature of Liebig's program was its ability to "breed chemists," to impart to developing researchers the art and the science of chemical research. If making high-fidelity copies of scientific practices is necessary for the distinctive rate of problem-solving we associate with modern scientific inquiry, then differences in the rate and degree of fidelity with which Liebig's lab was able to generate copies of scientific practice should partly explain the widely acknowledged qualitative difference between his lab's productivity and that of the rest of the world prior to that time.

2.1 Copying and Consensus

Practitioners in a given field vary significantly with respect to the practices they employ in the course of scientific inquiry. Thus it is natural to ask why some of these practices are imprinted on the next generation while others disappear from the population. In evolutionary biology, we think of a trait's probability of being represented in the next generation as a function of a number of factors, most importantly (1) demographic factors regarding various features of the distribution of traits in the population, (2) the existence and effectiveness of mechanisms of inheritance, and (3) ecological factors regarding how well suited a trait is for a given environment. The goal of this chapter is to describe the existence and effectiveness of some of the mechanisms by which scientific practices are often reproduced in the next generation. Before addressing that question, I can say something about how ecology and demography can shape the pool of practices that eventually exploits the available inheritance mechanisms.

According to the standard picture, individuals within a population can exhibit a range of variation with respect to which allele they have at a given locus. If there is no selection operating on that locus, the frequency distribution of alleles will randomly fluctuate from one generation to the next. If there is strong selection for a particular allele, all variation at that locus will be eliminated, leaving all individuals in subsequent generations with the

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

selectively advantageous allele; the allele has "gone to fixation." Selection will maintain the allele's fixation across successive generations for as long as it continues to be uniquely advantageous. It is possible for a deleterious allele to go to fixation in small populations, but not for it to be maintained across many generations in that way, because the absence of selection permits the introduction of new variations.

Scientific consensus is the scientific equivalent of allelic fixation (see Haufe ms-a). Following Kitcher (1993), let us call the group of practices that have gone to fixation "consensus practice." In the absence of selection, items on which there is some form of scientific consensus at one generation still have the highest probability of being represented in the next generation, because variation surrounding the issues on which there is consensus is very low. Due merely to inertia, apprentices will tend to inherit the practices of veterans. But due to the ongoing potential for new practices to emerge, scientific consensus can only be achieved and maintained through selection of one form or another.[3] As a consequence of the selective maintenance of consensus practice, idealized epistemic communities will exhibit pressure at the pedagogical level to ensure the direct high-fidelity copying of practices around which there is consensus. As Mulkay observed,

> science education . . . has traditionally operated to produce intellectual conformity. In all well-established scientific disciplines participants assume that, over a wide area, the correct problems have been posed and the correct answers have been found. Consequently, the student of science is required to acquire, rather than to question, the existing body of knowledge. (Mulkay 1991, 54)

One interesting corollary of this result is that scientific communities that are good at producing consensus will thus be good at copying their practices; the more practices on which there is consensus, the more scientific practice in one generation is likely to resemble scientific practice in a subsequent generation. Because this sort of high-fidelity practice copying is central to what it takes for scientific knowledge to undergo directional refinement, it follows that scientific knowledge will be a good candidate for directional refinement if scientific knowledge is good at producing consensus. And since directional refinement is what eventually causes consensus, it follows that a scientific community that has previously been good at producing consensus will continue to be good at producing consensus. Thus, for evolutionary theoretic reasons, we should expect to see a tight correlation between (1) high-fidelity copying and the propensity to achieve

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

consensus and (2) the propensity to undergo directional refinement and the propensity to achieve consensus.

Below I describe the mechanisms by which a parent generation's consensus practice is acquired by an offspring generation. The discussion is divided into two sections, the first looking at the acquisition of research problems (section 2.2.1) and the second looking at the acquisition of problem-solving strategies (section 2.2.2).

2.2 Copying Parent Practices

Scientific communities often face the challenge of persisting long enough to develop solutions to the problems they have chosen for themselves. Part of this challenge is already met by the longevity of individual scientific careers, during which individual problems can often be raised and solved. But many problems outlast the scientists that work on them. And at any rate, scientific communities are concerned not only with the continuity of specific research problems, but also (and often more fundamentally) with the persistence of certain kinds of problems—in much the same way as the stability of an ecosystem depends more on the persistence of certain ecological kinds or on the maintenance of certain functional ecological roles. Thus, communities must find some way of persisting that relies primarily on progenation—that is, on the production of new entities that can carry out future iterations of efforts to solve the problems that partly compose their selective environments. We know that some scientific communities have succeeded in this, because we know that they have succeeded in developing an edifice of knowledge that, like certain technologies, lies "outside the individual learning ability of an organism" (Mesoudi et al. 2013, 195). This is achieved by the training and professionalization of new scientists.

2.2.1 Persistence of Research Problems

Because the problem agenda in the many sciences is set by professional researchers, the training and professionalization of graduate students and recent PhDs is of particular importance for understanding how the specific problems defining a research field are inherited. Before reaching the stage where they begin professionalization, however, budding scientists will often have already begun to develop a sense for the fundamental character of research problems through their undergraduate science curriculum. This

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

consists of an acquaintance with, for example, the problems of motion, or of fields—general families of problems, for which the specific research problems they encounter are variations on an underlying unity. Developing a sense for what research problems look like in a given field is often crucial for the future professional scientist's ability to recognize and independently formulate novel research problems that fit naturally within the field. By transmitting the general form and substance of research problems in a given science, then, the undergraduate curriculum can help to lay a foundation for the inheritance of specific research problems and for the development of new, closely related problems. In this way, and for reasons that will later be made clear, the ontogeny of the scientist often parallels that of the organism; very general characteristics are acquired first, which then serve as a scaffold on which more specific structures can be elaborated.

One important mechanism by which research problems are inherited is—indirectly—through the more specialized formal and experimental training that graduate students and postdocs receive. With respect to formal training, as with biological morphology, the more specialized the formal technique is, the more narrow its range of application will be. Students trained in advanced formal approaches face limited opportunities for using them in a professional scientific setting. The opportunities that do exist are a combination of problems for which those advanced formal approaches were initially designed and problems that were developed subsequently by later generations of researchers. Thus, in the standard case, the continued use of those advanced approaches requires and, to an extent, promotes the continued relevance of the sorts of problems to which those approaches have historically been applied.

A similar phenomenon occurs in experimental training, where graduate students and postdocs typically acquire extensive experience in the use of particular technologies or particular model systems. As with formal techniques, advanced training in specific experimental techniques or on, say, a model organism like *D. melanogaster*, equips the budding scientist with a lot of highly specialized knowledge that cannot be expected to export easily to other, previously unrelated research problems. Due partially to this reason, partially to personal interest, a new scientist's independent research program routinely takes up the familiar sort of problems that she was trained to solve in her graduate or postdoc lab. In this way, in both formal and experimental training, the beneficial effect of increased specialization on

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

productivity can generate a phenomenon in science that is analogous to natural selection for habitat preference. In both scientific and biological contexts, a specialized set of problem-solving strategies is of greatest benefit to the agent (scientist or nonscientist organism) when those strategies can be applied to the problems they were designed to solve. Thus, in idealized contexts, we should expect to see selective pressure to replicate the challenges of the previous generation.

As effective as specialized training is for transmitting the research problems, it is relatively weak in comparison with the more direct mechanism of simply assigning research problems to graduate students and postdocs. Humanists like myself are often surprised to learn how little direct control graduate students in the natural sciences have over which scientific problems will occupy their doctoral research.[4] In the humanities and social sciences, a graduate student's long, circuitous journey toward his dissertation topic is littered with meandering, false starts, and self-doubt, all of which amount in the end to a topic that may or may not address a question of genuine importance. The ritual of choosing a dissertation topic is typically treated in these disciplines as a rite of passage, a sacred honor that defines the student as a researcher and a future member of a professional community. A great deal of emphasis is placed on novelty of problems and of methods, with a focus on distancing oneself from all existing work in the field (which, frankly, I find peculiar).

In stark contrast to this modern humanist tradition stand many natural sciences, where a graduate student's choice of dissertation topic is typically tightly constrained. Students entering a graduate program are either assigned to a lab or are allowed to choose the lab in which they will work. These labs are run by professional scientists with ongoing research programs that regularly demand significant amounts of labor. Once a part of the lab, the student is assigned a problem that fits into the managing scientist's existing research program. Even if students bring in external funding, they might still not be allowed to develop their own research problem. External funding merely affords them the freedom to choose from among a predetermined list of existing research problems.[5] Through expert peer review, both the research problems and the methods have already been vetted for propriety by the broader scientific community and have come to be established as components of the intellectual ecology and problem-solving strategy set, respectively. By limiting graduate students' choice of

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

thesis topic strictly to those research programs that are already underway, the scientific community effectively stabilizes the set of problems to which it is adapting over multiple generations. This is the most direct means by which research problems (or research problem types) are transmitted from veterans to apprentices.

2.2.2 Persistence of Problem-Solving Strategies

The same channels through which the research problems are inherited also support the high-fidelity copying of problem-solving strategies. Among the most important is undergraduate training. As with copying research problems, often a major part of what gets copied consists of an intuitive sense of the *primacy* of solving problems, as well as a sense of the *propriety* of certain approaches for solving certain problems. Through classroom instruction, textbook presentation, and lab exercises, students become acquainted with scientific inquiry as a ruthlessly task-oriented endeavor, decomposable into certain kinds of problems and the proper way to solve them. In this way, students of science come to place tremendous value on the acquisition and effective use of problem-solving strategies.

It is useful to divide the set of problem-solving strategies inherited through undergraduate education into two subcategories: (1) formal techniques and (2) experimental techniques. Kuhn observed in *Structure* and in his "Second Thoughts on Paradigms" (1977) that the use of problem sets in textbooks is the primary means by which students of the sciences are trained in the use of formal techniques. There he noted that the purpose of these exercises was to train students to see certain kinds of problems as the kinds of problems to which you apply some particular problem-solving strategy. For example, the introductory physics student will, through repetition, eventually come to see the problems of motion as being solvable through vector decomposition; after that, says Kuhn, "only manipulative difficulties remain" (305). The students who excel are those who are instantly able to identify a given problem on an exam as "the kind of problem for which you use thus-and-such" formal approach. This is an effective means of ensuring that the only students who proceed to more advanced training are those who have most successfully internalized the set of problem-solving strategies that are central to problem-solving success within a given discipline.

The transmission of experimental techniques typically proceeds in a much more experiential fashion than that of formal approaches (Collins

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

1992). Considered in the abstract, there is no obvious reason that the mode of transmission of the two approaches should differ. After all, experimental procedures are routinely described in scientific articles, on the principle that the description should be sufficiently rich as to allow anyone who reads it to be able to repeat the experiment. But things are rather more complicated when it comes to actually doing the experiments, the successful execution of which often requires the sort of experiential "tacit" knowledge that is characteristically high-bandwidth, difficult to articulate, and consequently much more faithfully transmitted through a master-apprentice sort of relationship.[6] It is a simple matter of fact that experimental proficiency cannot be achieved through the purely verbal channels, which often suffice in the case of formal techniques.

One dimension of undergraduate training that stands out to the philosopher and historian of science, as it did for Kuhn, is its nearly ahistorical and otherwise uncritical nature, and this should not be left unexplained. Kuhn provocatively described this feature of scientific education as pedagogically "unexceptionable," and suggested that to do otherwise could actually end up producing less-capable scientists (Kuhn 1962b, 140). Regrettably, Kuhn was probably right about this, at least as far as the practice of normal science is concerned. This is due to the fact that ignoring the historical and epistemological background to problem-solving strategies probably promotes their high-fidelity transmission to developing scientists.

The introductory physics student now proficient in the use of trigonometry for solving basic problems of mechanics will more than likely not be able to answer simple questions like "How is trigonometry able to tell us about force?" or "What is it about force that allows us to learn about it through the use of trigonometry?" He will not know anything about the challenges to inquiry that fueled the seventeenth-century search for an analytical system that could simultaneously depict magnitude and direction, nor even anything about the purely geometric derivation of the inverse-square law of gravitation.[7] But he can nevertheless identify and solve problems that would have been a major challenge for leading physicists of Newton's generation. Over the course of a semester, our student has been learning nothing but how to respond to different kinds of problems. The optimal result of this learning process is his ability to intuitively associate certain kinds of problems with certain kinds of analytical procedures. Knowing why we associate *these* kinds of problems with *these* procedures

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

does not necessarily improve his ability to generate the correct associations. Indeed, it may hinder his ability to do so, since he may insist on seeing the wisdom of doing things that way. Similarly, learning through historical and philosophical reflection that there is a variety of non-trigonometric approaches that are capable of treating the same problems may just as easily hinder his ability to add vectors, because it encourages the consideration of alternative problem-solving strategies and the reconsideration of the dominant approach's prestige.

These are suboptimal states, if our ultimate goal is getting him to reflexively—not critically—respond to certain kinds of problems in certain ways. (One thing I learned early on as a parent is this: Things work much more efficiently when the children learn to respond to my commands, rather than my reasons for giving them, because it is much easier to doubt a reason than a command.) In Kuhn's view, training scientists was fundamentally about preserving dogma, because dogma is what makes possible the unparalleled efficiency of normal science. Like any other dogma, scientific problem-solving strategies are most faithfully internalized when the recipients are given as little opportunity as possible to critically evaluate and contextualize them. To be clear, I am not advocating the suspension of critical thought in science per se; it is indisputably important for some aspects of scientific inquiry. But as a mechanism for imparting problem-solving strategies to the next generation, the foreclosure of opportunities for critical reflection does indeed seem "unexceptionable."[8] In the same way that paradigmatically Darwinian populations copy their traits with near perfection, idealized epistemic communities will be nearly perfect in their transmission of practice from veterans to apprentices. Ceteris paribus, we should expect a scientific community's propensity to undergo directional refinement (vertical progress) to vary directly with the fidelity of practice copying. Kuhn's physics-centric view arguably confirms this relationship. In the next section, we look at an exemplary episode from the history of science.

2.3 Breeding Chemists

If the growth of scientific knowledge is facilitated by the preservation of scientific practice across generations of researchers, research communities that excel at copying scientific practice with high fidelity should distinguish themselves from those who do not with respect to their rate of growth. A

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

clear illustration of this principle comes from the history of organic chemistry, which began to flourish in the 1830s under Justus Liebig. In 1824, Liebig took up a professorship at Giessen, a small university devoted to training civil servants and of no consequence in the world of scientific research. In less than two decades, Liebig would transform chemistry at Giessen from its initial state as a training program for a handful of future pharmacists into the most productive scientific laboratory in the world and the undisputed leader in research on organic chemistry. The section heading, borrowed from a classic (1972) study by historian Jack Morrell, explains how this remarkable transformation occurred: Liebig had developed at Giessen a system whereby productive research chemists could be systematically bred and exploited. As suggested by our model above, what Liebig had succeeded in doing, and what other scientific communities had been unable to do, was to develop an infrastructure that facilitated the high-fidelity copying of the problems and methods that were central to the discipline. It is for this reason that Liebig's laboratory is recognized as the first modern research group, and the exemplar on which subsequent university research across the natural sciences continues to model itself.

2.3.1 Liebig: Copying Problem-Solving Strategies

Liebig began his research career at an opportune time. Armed with the newly coronated atomic theory and a powerful heuristic for conceptualizing relationships between elements, chemists like Liebig could now gaze out on the vast array of substances as so much territory to be conquered. All that remained was to decompose substances into their constituent elements and compute their respective atomic formulas. The assignment of these formulas would be the focus of Liebig's research.

In order to assign a chemical formula to a given compound, a researcher must first determine the percentage by weight of every element in the compound. The standard approach to this task, in which Liebig himself had been trained, involved measuring the weight of different gases produced by a combustion reaction (Usselman et al. 2005, 2). An apparatus for analyzing organic compounds crafted by Liebig in collaboration with leading Parisian chemist Joseph Louis Gay-Lussac, for example, consisted of using copper oxide to oxidize the carbon and hydrogen produced through combustion into carbon dioxide and water. The percentage of hydrogen was calculated by weighing a sample of anhydrous calcium chloride before and after it had

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

absorbed the water, and then computing the fraction of the weight change due to hydrogen based on the percentage of water's weight known to be attributable to hydrogen alone. This procedure allowed for exceptionally accurate measurements of percentage by weight of hydrogen in an organic compound.

Circumstances related to the measurement of carbon were less happy. Unlike the more or less direct measurement of percentage weight due to hydrogen, carbon's contribution had to be determined volumetrically. In the Gay-Lussac–Liebig method, the carbon dioxide produced through oxidation was collected in a pneumatic trough, typically filled with mercury because of the relatively low tendency of gases to dissolve in it. The change in the level of mercury indicated the volume of carbon dioxide present, from which one could compute the fractional volume and weight due solely to carbon.

For many reasons, the volumetric determination of an element's percentage by weight is not very reliable, as indicated by the widely divergent estimates published by different chemists at the time (Holmes 1989, 143).[9] First, the fact that one was collecting gases meant the organic sample being analyzed had to be very small—typically only a few hundredths of a gram—because anything larger would have produced too much gas to be contained in the pneumatic trough (138). Such small samples did not make for reliable estimates. Second, one had to prevent the evolved gases from mixing with other gases, and there was essentially no way to guarantee this. Consequently, the mercury displacement could not be attributed with complete confidence solely to the evolved gas. Third, precision measurement of changes in volume required carefully controlled temperatures—both during and across trials—because any temperature difference would have an effect on volume, over and above the contribution made by the specific quantity of gas being collected. In addition to the unreliability inherent in the use of small samples, the need to meticulously attend to external conditions made precise estimates of carbon's presence extremely challenging to nonexperts.

The fulcrum about which Liebig's transformation of organic chemistry turned was his invention of a direct, gravimetric method for measuring the percentage by weight of carbon in a given compound, precisely analogous to the absorption method described above for the determination of hydrogen content. Rather than having to collect the carbon dioxide over mercury in order to measure its effect on volume, Liebig moved it through a series of bulbs, where it was absorbed by potash (potassium hydroxide), weighing

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

the potash before and after absorption. He was then able to calculate the percentage contribution of carbon to the analyzed compound.

The significance of Liebig's method is easy to appreciate when one compares it to the previous, volumetric attempts. In contrast to the sample size–limiting factor of collecting gas, Liebig's method allowed for the use of samples of virtually unlimited size; one merely needed to ensure that he had enough potash to completely absorb the evolved gas. In addition, because the potash would react only with the carbon dioxide, and because the device gave various indications as to whether a given trial had been air-tight, the threat of unwittingly measuring volume change due to multiple gases was greatly reduced (Usselman et al. 2005, 20). Finally, because the mass of carbon present in the compound was not sensitive to temperature, one did not need to carefully monitor external conditions as one must with volumetric methods.

One last distinction of Liebig's method concerned the plentiful signals that a particular trial was not poised to generate usable results. In their reproduction, Usselman et al. found that the Liebig process was replete with information that could instantly and easily allow a novice researcher to distinguish a good trial from a bad one:

> The manner of flow of product bubbles, the techniques of leak detection, and the taste of a bad analysis are some examples. . . . The timing of the first appearance of gas bubbles in the lye solution; the rate of evolution of the bubbles, and changes in that rate; the size of the bubbles, and changes in size; the timing and amount of rise in fluid level in the last of the three horizontal bulbs, indicating positive pressure of hot gases in the tube; retreat of the lye up into the first bulb at the end of the process, as the gases cooled; the stability of this elevated fluid level, indicating that the train remained air-tight; the absence of condensation of water vapour in front of the calcium chloride tube; the colour of the copper oxide over the course of the combustion—all of these, and others besides, provided many indications to the operator of a successful combustion. Some of these phenomena provided true feedback mechanisms that allowed the operator to optimize the process, such as continually adjusting the heating with a practiced eye on the bubbles in the bulbs, in order to maintain an optimally smooth and gentle combustion rate. Others were simple indications of failure: for example, if condensation appeared at the end of the combustion tube, if large gas bubbles passed completely through the kaliapparat, if the fluid level gradually fell after the tube was cool, or if the copper oxide bumped out of the combustion tube. In all of these cases the data would need to be thrown out, and the analysis repeated. (Usselman et al. 2005, 43)

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

The ability to easily detect high-quality analysis, coupled with the ability to easily adjust parameters in order to improve analysis quality, increased the precision, accuracy, and uniformity of results.

The critical point for my argument here is that both the method's simplicity and the ease with which good data could be sorted from bad made it possible for Liebig's approach to be easily copied. This is what is sometimes referred to in cultural evolution as a "content bias" (Mesoudi et al. 2013, 201). As Rocke (2003, 98) observes, "it was sufficiently simple and precise that analyses could for the first time be placed in the hands of relatively untutored students." A corollary to this was the fact that successful use of the method did not discriminate between "brilliant and average students. . . . It must be stressed that before Liebig perfected his 'combustion apparatus' the difficulties facing organic analysts were so formidable that only a few of the best men, such as [Joseph Louis] Gay-Lussac, [Louis] Thenard, [William] Prout, and [Jöns Jacob] Berzelius, had overcome them and produced accurate results before 1830" (Morrell 1972, 26). No longer was it the case that good measurements could only be obtained by a small number of virtuosi and their erstwhile apprentices. Most importantly, because "novices did not need personal instructions from a practiced chemist," "many chemists in the 1830s and 1840s successfully learned to perform the new method by simply reading about it." As a result, "within two years of its invention, the *Kaliapparat* had spread through nearly all the active centers of organic-chemical research, and had nearly completely displaced the earlier method" (Rocke 2003, 98–99).

The rapid diffusion of Liebig's approach to organic analysis is compelling evidence in favor of its revolutionary degree of reliability and simplicity. But a simple means of achieving reliable results was not the only mechanism by which Liebig's lab produced high-fidelity copies of a professional-grade research chemist. Morrell, for example, describes a training regimen developed by Liebig during which "each student was required to produce pure substances in good yield from raw materials." As every chemist knew, the reliability of an analysis was only as good as the purity of the substances involved; any student hoping to have a productive research career would need to master this skill. Once mastery had been obtained, the simplicity of Liebig's method enabled the student to "pursue original research and . . . work on his own under Liebig's general supervision" (Morrell 1972, 18). Effective training in how to prepare pure samples, combined with nearly

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

flat learning curve that characterized the use of Liebig's method, made it possible for a student to rapidly move from novice to competent researcher in fairly short order.

In addition to practical training, Holmes (1989, 128) observes that Liebig also instituted the practice—uncharacteristic among chemists at the time—of giving frequent exams in order to ensure that students were in fact acquiring the requisite knowledge and skill set. This is yet another means by which Liebig endeavored to promote fidelity in the chemist-copying process, having found that laboratory training alone "does not suffice, by far, to create in a young man some proficiency and skill in analytical work" (quoted in Holmes 1989, 127). By breaking the ontogeny of a chemist into well-defined stages, and by developing formal procedures for safeguarding the fidelity of the ontogenic process, Liebig could further guarantee that graduates of his lab would have obtained proficiency in each relevant aspect of research practice.

Morrell offers one further feature of the environment in Liebig's lab that helped to further promote high-fidelity copying: Liebig, the man. Morrell argued that Liebig's charisma inspired deep loyalty in his students, encouraging them to both emulate him and faithfully follow his instructions (Morrell 1972, 36). This points to another dimension of the copying process that evades descriptions of research pedagogy. Following Liebig was not something his students did in order to directly improve their proficiency as researchers. Nor was it something they did in order to increase the fidelity of the copying process. Or, in any case, following Liebig needn't have been something they did expressly for epistemic purposes in order for it to have direct positive effects on the ability of the Liebig lab to solve problems. In general, any mechanism that increases rewards for copying per se will tend increase the rate of problem solving, simply because copying is so central to the tendency of scientific practice to evolve adaptively. The ability to charismatically generate a stable of devotees is one way to achieve this effect.

Let us briefly review the circumstances by which Liebig was able to train relatively large numbers of students to reliably analyze organic compounds. The most significant factor was his invention of a method for measuring carbon content that required no expertise, no apprenticing, and very little general experimental training in order to produce reliable estimates. In a certain way, the Kaliapparat leveled the experimental playing field, collapsing the distinction between talented and untalented, novice and virtuoso, with respect to this centrally important analytical task. Second was

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

his decomposition of analytic training into separate components, each designed to ensure that a specific aspect of analytical skill had been successfully acquired. This decomposition would have allowed Liebig to identify more quickly and more specifically whatever was preventing a student from developing into a competent analyst who could work independently without direct supervision. Finally, we have Liebig's charismatic ability to inspire loyalty, which would have had the effect of increasing the diligence and effort with which his students approached the training process. Each of these elements had an additive effect on the rate at which Liebig was able to "breed" chemists, which in turn increased the rate at which Liebig's lab could solve problems in organic analysis.

2.3.2 Liebig: Reproducing Research Problems

Producing high-fidelity copies of competent methodologists was not the only means by which Liebig was able to dramatically accelerate the pace of research in his lab. Nor, in a sense, could it have been. According to the evolutionary model, the ability to propagate research *problems* across generations is equally essential to the adaptive evolution of scientific knowledge over time.

The inheritance mechanism by which research problems were copied in Liebig's lab is familiar to anyone acquainted with graduate training in the natural sciences. Liebig would simply specify the problems on which his students were to work (which was itself possible only because of the relative ease with which he was able to train them to do independent work). These assignments can be divided instructively into two types. The first kind of problem was essentially a "gap-filling exercise," such as "an individual problem arising from a secondary question left unanswered by one of his own investigations, an unresolved question raised in the recent chemical literature, [or] an analysis made by another chemist that Liebig thought needed confirmation or a critical reexamination" (Holmes 1989, 155).

These problems would have been structurally similar to ones that Liebig himself had already solved, as well as to those lab exercises on which the student had previously been trained. The second sort of assignment—which, Holmes argued, emerged somewhat later—was one in which Liebig would instruct an advanced student to work on a qualitatively novel project "clustered around a common problem closely connected with Liebig's own most immediate research concerns" (Holmes 1989, 155).

These two types of assignments contributed to the evolution of scientific knowledge under Liebig in different ways. The gap-filling exercises can be

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

seen as an attempt to maximize the benefits that can be achieved through making perfect copies of existing practices—no new methods, no new types of problems, just making the absolute most of an existing set of practices. Liebig's habit of assigning qualitatively novel projects, on the other hand, was an attempt to increase the kinds of benefits that scientific practice is capable of generating by introducing some directed variation into the set of problems that existing methods are able to handle. In ecological terms, the gap-filling assignments attempt to exhaust the resources exploitable by a given morphology in a given environment. The latter sort of assignment attempts to expand the environment in order to thereby expand the set of resources that are exploitable with a given morphology.

In the 1840s, the research problems that Liebig assigned to his advanced students began to bear less of a relation to his own research problems than they had during the period that defined the lab's meteoric rise in the research community (Holmes 1989, 163). At this point, we see a decline in the productivity and significance of the Liebig lab at Giessen. If we understand this decline as a symptom of the reduced pace of scientific progress, the evolutionary model explains why it coincided with the end of Liebig's practice of decomposing his research problems into smaller projects that he could then assign to students: Liebig was no longer generating copies of the research problems that defined his position on the frontier of organic chemistry. The sheer breadth and complexity of the problems required an enormous amount of human labor; there was simply no way around that. At the apex of Liebig's productivity, the labor was distributed across a reasonably large number of individuals, each of whom could be trusted to independently produce reliable results. The enviable rate of productivity was not lost on the French chemist Jean-Baptiste Dumas, who

> in 1837, 1840, 1846, and 1847 . . . wrote pleading official reports to the ministry of public instruction, deploring the lack of official laboratory facilities in Paris and often urging a model strikingly reminiscent of the Giessen facility and organization. It is not necessary, he wrote, "to wait for a question to be resolved by the individual work of one of the [the faculty's] professors extended over several years, when it can do so in a few weeks under his direction by the collective effort of a dozen beginners in science." (Rocke 2003, 112)

The necessity of copying problems within and across generations of researchers is not something that was confined to the community of research chemists during the mid-nineteenth century. Rather, I would argue that it is an inevitable by-product of empirical inquiry. The kinds of

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

problems that are represented in future iterations of the scientific community are the kinds that are good at copying themselves. One way in which a problem can be good at copying itself is if it is intrinsically interesting, in the way that many of the ancient problems of philosophy are. Some problems in nature are like this, but many are not. What makes empirical problems particularly good at copying themselves is not their general interestingness, but the fact that their study can quickly become extremely complicated and labor intensive. Once a problem passes a certain threshold level of complexity or labor intensity, no single individual is capable of completing all the work necessary for solving it. The problem must be decomposed into smaller, more individually manageable projects, each of which helps to form part of the overall problem's solution. This decomposition allows the problem to then be distributed across several individuals. The historic shift in the way scientific research was organized occurred as a necessary by-product of Liebig's uptake of problems that were sufficiently labor-intensive (gap-filling exercises) and sufficiently complicated (problems at the frontier of scientific knowledge) to warrant their widespread distribution.

As a final point, consider the effects of high-fidelity practice copying on the scientific community's ability to achieve consensus, and the associated implications for future inquiry. One of the central points made by Kuhn in *Structure* is that the significance of the transition to paradigm-governed science lies primarily in how it allows practitioners to *take for granted* certain ideas about nature and its proper study (Kuhn 1962b, especially 10–22). The permission to take certain items for granted derives from their widespread or near-universal acceptance. Once there is widespread agreement concerning some matter of fact, theoretical representation, or problem-solving approach, practitioners are no longer obligated by the epistemic norms of scientific practice to justify the advance to more refined stages of inquiry that take that item as given. The ability to take certain things for granted thus enables practitioners to develop inquiry beyond the most rudimentary stages, to deepen knowledge in the manner symptomatic of populations dominated by the influence of selection. Kuhn's image of periods of normal science reflects a community that, through its adoption of an almost universally accepted paradigm, becomes progressively better at adding to the pile of things that can be taken for granted.

In his study of Liebig, Holmes argues that the Kaliapparat's widespread diffusion strongly accelerated the pace at which consensus on organic composition was reached. Because the Kaliapparat made reliable results so easy

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

to obtain, and because everyone was using it, chemists working independently were more or less guaranteed to achieve similar estimates of carbon content. Although Holmes (1989, 142) is in part correct in stating that "movement toward a standardization of the basic analytical procedures narrowed the area for disagreement," standardization alone is not sufficient to explain the narrowing of disagreement in this episode. For instance, a standardized method that was difficult to use would be less likely to narrow disagreement among practitioners, because different people are more likely to get different results. It was the ease with which proficient Kaliapparat use could be copied that led to both standardization and consensus on organic composition.

2.4 From Consensus to Refinement

I have been arguing that the epistemic consequences of high-fidelity copying of scientific practices possess the same evolutionary significance as do the adaptive consequences of high-fidelity copying of biological traits: copying is necessary for the process of directional refinement that characterizes progress in both contexts. I have illustrated the special epistemic consequences of high-fidelity copying by drawing attention to the effectiveness with which scientific practice was copied across—and eventually beyond—Justus Liebig's lab. I have suggested, following the consensus view among historians of chemistry, that the successful copying mechanisms associated with Liebig's lab were chiefly responsible for that lab's historically unprecedented success.

However, I have yet to say much about what that success was like, apart from the fact that there was a general uptick in the rate at which the research community was able to agree on organic formulas. That does not give us much on which to ground the assertion that high-fidelity copying is necessary for the directional refinement of scientific practice. I now want to draw a more substantive and direct causal connection between the success of copying in Liebig's lab and the directional refinement of a specific aspect of scientific practice in nineteenth-century chemistry: the development of structural chemistry. I will base the argument for this connection on three claims: (1) that the development of structural chemistry was at its core the refinement of the set of questions chemists viewed as relevant to the understanding of a substance's nature; (2) that there was something about

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

this development that required high-fidelity copying of scientific practice; and (3) that the specific copying mechanisms associated with Liebig's lab were particularly well suited for the sort of practice-copying that led to the refinement of chemists' questions. We already established the case for (3) in section 2.3.1: the Kaliapparat was just spectacularly easy to use and, for that reason, it diffused rapidly across the research community. Let us now see whether there is support for claims (1) and (2).

The dominant view among chemists prior to 1826 had been that each substance is associated with a unique chemical formula. As Rocke (2003, 93) evocatively describes the situation, "The idea that two distinct substances could have the same chemical composition probably must have seemed as unlikely to them as, in this century, the possibility that two unrelated people might have identical fingerprints or even identical genetic material." Thus, to understand why a substance was the way that it was, a chemist need only ask two questions: (1) "What are its basic constituents?" and (2) "How many atoms of each constituent are there?" But in 1826, a highly publicized debate between Liebig and fellow German Friedrich Wöhler led to the undeniable realization that the answers to both of these questions were the same for two distinct substances, fulminic acid and cyanic acid (Rocke 1984, 171–174). As analytical methods continued to improve, more identical formulas were uncovered, and by 1830 Berzelius had elevated the pattern to the status of a general phenomenon, which he called "isomerism."

The discovery of isomerism meant that knowing the kinds and numbers of atoms composing a substance was not going to be sufficient for explaining its properties. Since the advent of Daltonian atomism at the beginning of the nineteenth century, chemists had been toying with the idea that the physical arrangement of a substance's atoms might influence its nature. These speculations do not seem to have been responses to theoretical or empirical challenges, such as those raised by isomerism. Rather, they were a by-product of the spread of the mechanistic philosophy within chemistry, which began to accelerate during that period. The mechanistic approach to explanation highlights a precious few quantities as being relevant to understanding an entity's behavior, one of which is the arrangement of its parts. Thus, the potential chemical significance of atomic structure would have been (and was) widely acknowledged before the formal announcement of isomerism as a real phenomenon (Rocke 1984, 167–169). But as chemists began to appreciate the generality of isomerism, appeals to structural

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

arrangement moved from the realm of idle speculation and into the arena of essential explanatory posit. By 1830, Berzelius had officially embraced the obvious refinement to chemists' battery of questions: knowledge of elemental composition must be supplemented by information concerning "the relative positions of the elementary atoms."[10]

This, then, is the argument in support of the notion that the advent of structural chemistry amounted essentially to a directional refinement in the set of questions chemists took to be explanatorily relevant. The likely influence of structure on a substance's behavior was one of a few very general and direct consequences of mechanism, one that had been widely appreciated among chemists immediately following the embrace of mechanism. Thus, we would expect them to appeal to structural properties in the event that they were faced with behavior that they could not explain via reference to elemental composition. The opportunity to make that appeal appeared with the discovery of isomerism. Prior to that point, there was no reason to think that kind and number of atoms would not have been explanatorily sufficient.

The second claim is that the development of structural chemistry as an active area of research required the high-fidelity copying of relevant scientific practice. The task of uncovering exactly how differences in the atomic arrangement of isomers explained the differences in their chemical character essentially required two types of laboratory-intensive work: (1) the analysis of organic compounds, where "isomers were known to proliferate,"[11] and (2) the search for isomers, once an organic compound had been analyzed. As Dumas indicated in the passage quoted in section 2.3.2, there are only two ways to do this: either (1) one researcher working over many years or (2) many researchers working over a relatively few years.

Concerning analysis, the Wöhler–Liebig dispute over fulminic acid and cyanic acid made clear that isomers posed a special challenge for the reliability of formula assignment that was much easier to resolve at the community level. If a single researcher produced identical formulas for two distinct substances, the most probable explanation for the identity would have been measurement error. Some mistake—contamination, sample impurity, sample size—would have been a much better candidate than the purely speculative idea that the formulas genuinely are the same, despite the fact that the substances are so different. Researchers must also have been aware of the very high probability that their approach simply may

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

not have been sensitive enough to detect differences in numbers of atoms. Whatever the reason, one always had to harbor the suspicion that identical measurements of an element's content did not necessarily reflect identical numbers of its atoms. The suspicion is even greater when whole formulas turn out to be identical, because of the background principle, which stated that what distinguishes substances from one another is their elemental composition only.

The candidacy of measurement error as an explanation for formula identity was significantly weakened by the fact that different researchers in different labs, procuring their samples from different sources, were producing identical formulas on a routine basis. By varying as many aspects of measurement as possible, researchers could be reasonably confident that inter-researcher agreement reflected underlying facts about the substances themselves, rather than the common weaknesses of the researchers. Under these new conditions, members of the research community had a compelling empirical reason to question their commitment to the "one substance, one formula" principle. The ability to adjudicate these disputes at the community level was made possible only through high-fidelity copying of scientific practice in this case (this is essentially the point made by Holmes regarding the "narrowing of disagreement").

Once a compound had been analyzed, chemists were in a position to start making guesses regarding how its atoms were arranged. The only way to really start learning anything about arrangements was to develop models of how the atoms in a given molecule might possibly be arranged and then derive predictions for how many isomers a given arrangement could possibly accommodate. For example, if the atoms were linked together in a chain, they might bond in either of two possible ways; whereas, if they formed a ring, the same group of atoms might bond together in any of three possible ways. The number of possibilities that a given structural arrangement can accommodate corresponds to the number of isomers we expect that formula to have. In this way, different models of the relative positions of the atoms for a particular formula gave rise to different predictions about how many isomers we would expect to be associated with that formula. As Rocke (2001, 416) argued, "Positive evidence in chemistry is also often examined only by massive testing. The structure theory itself generated literally infinite numbers of 'small' predictions about what compounds ought to exist and what compounds might well be impossible to prepare."

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Suppose that we have one structural model that predicts two isomers and another structural model that predicts three isomers. If only two isomers have been found, we face the question of whether we have found all of the isomers or only two of them. The only way to distinguish between the two situations is to keep looking until we're confident that we would have found a third isomer, had one existed. And the only way to be confident of that is to test a great many possibilities.

The development of research into the structural arrangement of atoms in a compound was a directional refinement in the set of questions with which chemists had hitherto been occupied. By "directional refinement," I mean the next step in a sequence of questions defined by the preexisting commitment to the mechanistic philosophy. That questions about structure are part of the framework of mechanism is not disputed, and there is plenty of evidence that chemists were contemplating structural questions at a very hypothetical level in the period immediately following the adoption of mechanism. The fact that they were contemplating these questions at such an early stage, without any explanatory need, strongly suggests that the structural consequences of the mechanistic philosophy rationally compelled them to suspect that structural properties would be causally influential in some way. But the actual investigation of structural influence had to wait until (1) there was some phenomenon that cried out for a structural explanation (e.g., isomerism) and (2) there was some way of empirically investigating structure. Confirming models of structural arrangement required that chemists be able to generate organic formulas and then see whether the elements could combine in the number of ways predicted by the models. The only way to do this is to just keep trying to combine the constituents in novel ways: there is no shortcut. Indeed, Rocke argued that the laboratory-intensive nature of this search process explains why Germany, so well equipped with labs and the students to populate them, was able to oversee the development of structural chemistry. By contrast, the relative paucity of lab space in France made it impossible to move forward: "Within 3 months of publishing his first brief article on benzene theory [a theory of the structure of benzene], Kekulé wrote to a friend, 'A great deal is in the works; the plans are unlimited, for the aromatic theory is an inexhaustible treasure trove. Now when German youths need dissertation topics, they will find plenty of them here'" (Rocke 2001, 417).

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

2.5 Conclusion

I began this chapter with the intent of presenting a detailed argument for the connection between the directional refinement of scientific knowledge through natural selection on the one hand and the ability to make high-fidelity copies of scientific practices on the other. The first part of the chapter involved the use of basic principles of evolutionary theory and ecology to elaborate on this connection. The second part of the chapter looked at the variety of channels that are built into the social fabric of many branches of scientific research and that promote the high-fidelity inheritance of scientific practice, conceived of as the sum of problem-solving strategies and the kinds of problems to which they are applied. The third and final part of the chapter looked at the historical phenomenon of Liebig's lab.

The relevance of Liebig's lab to the connection between directional refinement and high-fidelity copying is threefold. First, Liebig's lab is recognized as the first modern scientific research group, and the first group capable of solving empirical problems at the distinctive rate associated with modern scientific research. The group's unprecedented rate of problem-solving was acknowledged at the time. Second, the lab's distinctive rate of problem-solving was partly explained by the high-fidelity transmission of research problems and methods from Liebig to his students. Following the criteria for Darwinian populations, we can unambiguously assign a direct causal role to Liebig in explaining this transmission. Third, this ability to generate high-fidelity copies of chemistry's cutting-edge problems and methods was an essential ingredient in the chemical community's transition to the investigation of structure, the next phase of chemical inquiry as defined by the community's preexisting commitment to mechanism.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

3 Isolation of the Evolving Population

Even if scientific practice copies itself with high fidelity, high-fidelity copying cannot by itself establish scientific practice as a Darwinian population. We now turn to another dimension along which the candidacy of scientific practice as a Darwinian population must be assessed. In addition to establishing that copying is high fidelity, we also need to answer a central question about what accounts for the different rates at which different practices copy themselves—namely, to what degree are those different rates explained by intrinsic features of the practices themselves? Put another way, what portion of a practice's differential reproductive success is explained by its goodness of fit—its *adaptedness*—to its environment? This question suggests that it is possible for differences in adaptedness to a given environment to fail to fully account for differences in reproductive success. Indeed, there are a variety of factors that can weaken the contribution of adaptedness to actualized reproductive success. One of the more prominent factors is population size; ceteris paribus, the smaller the population is, the less reliable differences in adaptedness will be as a way of predicting what future compositions of the population are going to look like. So, small populations are poor candidates for being paradigmatic Darwinian populations, because differences in actualized reproductive success depend less than they otherwise might on differences in intrinsic features.

Were we to show that differences in the intrinsic features of practices explain differences in their reproductive success, we would have established that what matters to the rate at which practices copy themselves is how well they respond to the pressures raised by their environment. This is necessary to establish that the directional tendencies within a science are a consequence of natural selection. But more is required to show that that directionality is *epistemic*—that is, that the kind of historical pattern often

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

described as "scientific progress" ought to be interpreted as a trend in the historical development of scientific knowledge toward increasing degrees and kinds of specialization. To make this case, we must also see whether the pressures to which a community's practices respond—and which explain why some practices produce more copies than others—are the kinds of pressures that discriminate between practices on the basis of how much they add to scientific knowledge. If the pressures to which a community's practices respond instead discriminate on the grounds of, say, how politically expedient those practices are, then the directionality of science in that community is not a signal of epistemic growth.

In this chapter, I make two arguments, jointly aimed at supporting the thesis that epistemic pressures contribute more to changes in some communities' practices than does any other factor. I begin by arguing for a method by which we can gauge whether the most significant pressures exerted on scientific practices in certain communities are epistemic. We can determine this, I claim, by noting the prevalence of factors that isolate the population of scientific practices from the introduction of nonepistemic influences. A community's ability to isolate itself from nonepistemic influences is an index of its members' propensity to behave as epistemic niche constructors. The significance of epistemic niche construction for understanding the particularly powerful instances of problem-solving in modern science cannot be overstated. Ultimately, I argue that the adaptive power of niche construction in biology is of a piece with the epistemic power of the ability of certain scientific communities to stabilize their environment and isolate themselves from the introduction of deleterious variations in scientific practice. Ceteris paribus, the degree to which a scientific community isolates the population of practices from nonepistemic influences should vary directly with the degree to which directional changes in its practice are epistemically driven.

I then argue for a way of understanding whether the exertion of selective pressures—rather than, say, evolutionary drift—plays the strongest role in determining how the population of practices changes over time in a scientific community. This, I argue, can be inferred from what we can know about those pressures' magnitudes, and from what we can know about the tendency of scientific research problems to get solved. The combined result of these influences is a two-stage filtration process—isolation and selection—that determines the extent to which directional refinements to practice amount to epistemic improvements. A Darwinian population's

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

propensity to be refined by selection is partly a function of its degree of intrinsicality. Likewise, a scientific community's propensity to undergo vertical progress is a function of the degree to which modifications to practice are driven by internal pressures.

The chapter proceeds as follows: I begin in section 3.1 with an introduction to the concept of niche construction, followed by a discussion of some of the special features of niche construction that strongly promote directional evolution. I then describe in section 3.2 how scientific practices could occupy a carefully constructed niche composed of selective pressures that discriminately sample scientific practices for representation in future generations, increasing the susceptibility of scientific practice to directional refinement through selection.

Moving from here to the conclusion that epistemic factors are the dominant contributors to chance in scientific practice will require two additional steps. The first step, presented in section 3.3, is to describe conditions under which the pressures that define the constructed niche occupied by scientific practice would be epistemic. Scientific communities are characterized partly by particular sets of filters aimed at influencing the propagation of practices within the community. We can imagine how those filters might discriminate between practices on the basis of epistemic criteria. The second step involves providing reasons to think that these filters are generally efficacious in certain circumstances. This is the subject of section 3.4, where I turn to describing conditions under which selection pressures matter more to the propagation of scientific practices than does any other evolutionary factor (e.g., drift); call this the "Intrinsicality Thesis," in reference to Godfrey-Smith's *intrinsicality* condition on Darwinian populations. In particular, I emphasize the general evolutionary role of niche construction in buffering against the effects of mutation, migration, and drift, all of which on average oppose the incremental increases in adaptedness promoted by natural selection. Establishing this claim will conclude this portion of the argument for how to assess the susceptibility of scientific practice to modification through selection.

3.1 Niche Construction and Directionality

A number of factors affect a population's ability to undergo directional refinement. The ideal scenario for directional refinement of traits would be

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

one involving absolute stability in a small set of environmental factors that make a difference to reproductive success. Absolutely stable environmental conditions stabilize the phenotypic optimum, which is typically understood as the version of a trait that maximizes reproductive success. Under natural selection, a trait is favored in proportion to its relative proximity to that optimum compared with other traits in the population. As long as environmental conditions remain stable, the mean trait value will continue to evolve in the direction of the optimum.

Contrast this scenario with the opposite extreme—total environmental instability. Directional refinement under such conditions is impossible, because there is no phenotypic state that is consistently favored by the environment. At one moment, one trait value is favored; at another, the opposite value might be favored. Because no phenotypic state is consistently favored by the environment, there is no way for traits to increase their relative proximity to a favored state. Selection cannot push mean trait values to evolve in any particular direction, and we thus expect no increases in mean reproductive success.

We can think of these two extremes—total stability and total instability—as occupying two poles of a continuum, movement along which constitutes changes in the degree to which traits can profitably specialize in solving a given selective problem. Specialization pays big dividends under conditions of perfect stability, because the persistence of certain selective problems means that it is always beneficial to develop better solutions to them. Other things being equal, improving one's performance on a given task always increases reproductive success, so long as performing that task remains relevant to reproductive success. Under conditions of perfect stability, the tasks that are relevant to reproductive success do not change.

When the relevant tasks do change, continuing to improve one's performance on formerly relevant tasks is disadvantageous. More importantly, the rate of change affects the degree to which a population can increase its mean reproductive success. Perfect stability allows mean reproductive success to maximize, but all that is really required is that stability endure long enough.[1] If environmental change is relatively infrequent, a population can maximize mean reproductive success relative to a given set of selection pressures. Frequent environmental change, in contrast, generally impedes a population's ability to increase its mean reproductive success. In this way, mean reproductive success can be understood to vary directly with the

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

degree of stability exhibited by selective pressures: stability rewards specialization, and specialization increases reproductive success.

With the foregoing in mind, imagine two types in a population, *A*s and *B*s, who differ in their propensity to wander in space. If an *A*'s range is relatively smaller than a *B*'s, *B*s are exposed to a wider range of environmental conditions. And thus *B*s are likely to encounter more factors that affect reproductive success than are *A*s. Now, since the *A*s' selective environment is more stable than *B*s', *A*s are better able to increase mean reproductive success. There is a stable target to which *A*s can adapt, and so they are able to make incremental progress toward it, with improved performance causing greater reproductive success. By contrast, the wide-ranging *B*s find themselves pulled in relatively many directions at once and are consequently less able to improve their overall performance relative to their environmental conditions. Descendants of *B*s that happen to have more restricted geographic ranges over which they wander will have more offspring, because they are more likely to be successful in dealing with their relatively small set of selective demands than their parents were in dealing with a larger set. In this way, selection will restrict the range of *B*s to a relatively small, reproductively optimal patch.

Because of the general selective pressure to constrain the number of factors affecting reproductive success, we expect species to evolve strategies for restricting and stabilizing their selective environments. The most common of these is a tightly restricted geographic range, or what is called *endemism*. The vast majority of species on Earth are endemic, which in part accounts for the bewildering diversity of biological forms across time and space. The tendency of endemics to remain in a relatively confined geographic area limits the number of environmental phenomena they encounter, and thus limits the number of factors that might affect reproductive success.

The prevalence of endemism invites questions about what sorts of organisms tend *not* to be endemic, but rather geographically widespread. Unsurprisingly, many of these organisms have minimal resource requirements and can basically survive under any conditions (e.g., plankton). Others have fairly specific requirements that are themselves geographically widespread (e.g., planktotrophic organisms). Consistent with the adaptive logic of endemism, marine organisms tend to be more geographically widespread than terrestrial organisms, owing to the relative homogeneity of marine environments as compared with terrestrial ones.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Over the past 30 years, evolutionary biologists have begun to increase their focus on an important means by which some species are able to extend their geographic range across highly varied environmental conditions without exposing themselves to the kind of selection that restricts geographic range in order to maximize reproductive success. This phenomenon, known as *niche construction*, allows organisms to determine and recreate the environmental factors that significantly impact their reproductive success. Unlike endemism, where selection confines a species to a particular plot in space, niche constructors often have the ability to manufacture the conditions that matter to their reproductive success across a wide variety of environments. Nowhere is this more clear than in the case of human beings, who live in every terrestrial environment (also in outer space) by manufacturing the conditions relevant to survival and reproduction wherever they go.

Being able to manufacture the conditions that affect reproductive success has important consequences for a population's susceptibility to directional refinement through natural selection. One of these important consequences is the population's ability to stabilize a phenotypic optimum by stabilizing the set of factors that matter to survival and reproduction.

The evolutionary fate of organisms that are not niche constructors is strongly tied to changes in their external environmental conditions. For them, changes in those conditions constitute changes in the challenges faced by an organism. These changes, in turn, affect both (1) what constitutes an optimal phenotypic state and (2) the degree to which populations are able to realize that phenotypic state in evolutionary time. Niche construction, however, can result in the decoupling of a population's evolutionary fate from changes in its external environment. Because the population is in control of which sorts of things influence reproductive success, the evolutionary significance of environmental variation in time and space is diminished; populations can manufacture their selectively relevant conditions under a range of environmental circumstances.

The ability to manufacture their selectively relevant conditions across different environments allows populations to stabilize those conditions across generations, a phenomenon known as *ecological inheritance* (Odling-Smee 2007). With their adaptive demands decoupled from the external environment, niche constructors can transmit the selective relevance of certain factors to subsequent generations without needing to depend on

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

the stability of the external environment. The persistence of these factors across several generations allows the population to continue on the adaptive trajectory defined by the now-stable phenotypic optimum. As long as the population is able to stabilize the set of selection pressures, it will continue to refine traits in the direction of the associated optimum. While the phenotypic optimum remains stable, the best way for an organism to increase its fitness is to decrease its distance from that optimum.

To summarize, with niche constructors, selective pressures can bear a spatially arbitrary relationship to the external environment, allowing for the persistence of selection pressures across generations (ecological inheritance) under a wide range of environmental conditions. The persistence of these factors across generations stabilizes the phenotypic optimum, and stable phenotypic optima facilitate directional evolution.

A changing environment is not the only obstacle facing the population's prospects for optimization. Even when a phenotypic optimum has been allowed to stabilize, the population's ability to occupy it depends on what sorts of variation appear in each successive generation. The appearance of variants that lie closer to the optimum facilitates optimization, while the appearance of variants that lie further away impedes it.

A variety of factors influence which variants appear in a given generation. In chapter 2 we looked at the most significant influence on those variants, reproduction by the previous generation. If every member of the previous generation perfectly copies itself—and no other factors influence the population's makeup—the subsequent generation will be a perfect copy of the previous generation. That is, it will contain the same types at the same frequencies.

Of course, other factors are always involved. Notably, natural selection affects the makeup of the next generation by favoring the reproduction of some variants relative to others. By changing frequency distributions of variants in light of differences in adaptedness, natural selection gradually eliminates all but those variants that are best adapted to the selective pressures that characterize that population's selective environment.

The ability of selective elimination of the unfit to alter the population's composition depends crucially on the degree to which the population is isolated from the introduction of new variants. This is because new variants can, on average, be expected to be less well adapted than the prevailing mean for that population. Let us imagine a group of organisms,

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

split into two distinct populations, *A* and *B*, each of which is evolving in the direction of its respective local optimum. Given sufficient time apart, and sufficiently strong selection, the two populations will begin to diverge morphologically; the *A*s will continue to ebb closer to the optimum they previously shared, while the *B*s will redirect toward their distinctive new optimum. With each successive generation, *A*s and *B*s diverge further and further from their previously shared mean trait values, and thus further and further from each other.

As the distinction between their phenotypic means increases, so does the degree of difference between the variants that appear in an *A*-generation and the variants that appear in a *B*-generation. As each population continues to shore up the distance between its population mean and its local optimum, selection eliminates variants close to the original, general mean held prior to the split. Consequently, the average *A* variant is, in each successive generation, further from *B*'s local optimum—most importantly, the average *A* is further from the *B* optimum than the average *B* (and vice versa). The introduction of *A*s via migration will, in turn, disrupt the directional evolution of *B*s toward the *B* optimum, by slowing down or deflecting the *B*s as selection attempts to push them toward it. While selection is busy exterminating the unfit, migration reintroduces the unfit (i.e., with respect to the *B* optimum) back into the population.

It is thus critically important for the directional evolution of *B*s that the flow of *A*s into the breeding population be curtailed. Once the population is isolated from the migration of *A*s, the variants that appear in each successive *B*-generation will be limited to descendants of members of the previous generation of *B*s. If the ability of members of the previous generation were affected by how well adapted they were to their environment—that is, if natural selection were operating—well-adapted *B*s can be expected to have left more descendants than others. As long as the *B*s remain isolated, each iteration of this process brings the *B* phenotype closer to the *B* optimum.

If isolation promotes directionality, so too do phenomena that promote isolation. Niche construction is one such phenomenon. Because of their ability to arbitrarily establish the selective significance of some environmental factor, niche constructors are able to rapidly increase the magnitude of a selection pressure. When the relevant selection pressure affects the opportunity for certain kinds of individuals to reproduce, niche constructors can effectively

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

isolate themselves from the introduction of new variants into the reproductive pool. Preventing the introduction of new variants into the reproductive pool, in turn, allows the relevant population to continue undeterred toward its local optimum.

The family of phenomena associated with mate choice provides an intuitively graspable illustration of this process as it occurs in nature. The clearest example is mate preference, which possesses both the element of arbitrariness and functions to preclude certain individuals from the mating pool. Females in many species often display strong preferences for particular kinds of males—males that make a particular kind of song or have a certain kind of plumage. Depending on the degree of intensity with which these preferences are exercised, they often prevent all but certain kinds of males from reproducing. In effect, these mating preferences isolate the population from the introduction of new male variants, allowing evolution to proceed in the direction dictated by the mating preference. In some cases, mate preferences function in such a way as to consistently favor deviations from the mean trait value. Under these conditions, preferences can generate a positive feedback loop that results in the highly exaggerated ornamental displays for which many males are famous (e.g., peacocks, bowerbirds)—a process known as *Fisherian runaway*. The tendency of mating preferences to isolate the mating pool and promote directional divergence is widely understood as being one means by which new species can arise.[2] Such is the directional power of mate preferences.

Of related importance is the fact that mate preferences are often credited to contingent features of neural architecture. Females might prefer songs or croaks of a certain frequency, plumage of a certain hue, or dances of a particular style—for no reason other than the fact that their sensory mechanisms are arbitrarily biased in that direction. That is, these preferences need not function so as to detect males with certain environmentally advantageous qualities (although they certainly can evolve in this way). Rather, the simple fact that females are sensorily biased in a certain way can be enough to cause the directional evolution of certain traits in males. Those traits need not be adaptive in other ways, nor need they signal the presence of traits that are adaptive in other ways.

Apart from its utility as an illustration of the causes and evolutionary consequences of reproductive isolation in nature, mate preferences are especially useful as a way of bridging any perceived gap between the

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

evolutionary process in biological contexts on the one hand and scientific (or other cultural) contexts on the other. Just as a mating pool can be, to a significant degree, determined by the arbitrary preferences of existing members of the population, so too can the pool of scientific practices be merely a consequence of the preferences of practicing scientists. Viewed in this way, the population of scientific practices can be understood as occupying a constructed niche, where the selection pressures are the preferences of scientists. Like female mate preferences, scientists' preferences determine which scientific practices gain entry into the reproductive pool, and how long those practices are allowed to stay. It is largely for this reason that the empirical question of what scientists' preferences actually are is so significant for differentiating between epistemic and sociocultural directionality.

In section 3.2 I describe conditions under which scientific practices can be conceived as occupying a niche constructed by the epistemic preferences of scientists. In idealized epistemic communities, the niche is constructed purely out of epistemic preferences of scientists, and the modification of scientific practice is dominated by refinement in the direction of the optimal fit with those preferences. Actual scientific practice can be expected to undergo directional epistemic modification in direct connection with the degree to which a given community's niche is constructed out of epistemic preferences.

3.2 Scientific Practices in a Constructed Niche

In section 3.1 we established three points that bear significantly on the question of whether scientific practices are good candidates for directional refinement. First, we noted that stabilizing the set of factors that affect reproductive success facilitates directional refinement. Second, we saw that organisms that are able to manufacture the factors that affect their reproductive success—niche constructors—are good at stabilizing those factors, because they are able to maintain their selective significance across a broad range of external environmental conditions. Third, we saw that certain kinds of selection pressures, like mate preferences, can exclude certain variants from the reproductive pool, which isolates the population from the introduction of migratory variation. Like stabilizing selective factors, isolation also promotes directional refinement in that it preserves the adaptive gains made by natural selection in previous generations and allows the current generation to build on those gains. It is important to note that, for niche constructors who inherit their ecology from the previous generation,

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

isolation from the introduction of new variants helps to prevent the introduction of new selective pressures.

I now want to illustrate some of the ways in which preferences of practicing scientists contribute to the construction of the niche occupied by scientific practices. These preferences help to isolate scientific practice from the introduction of new variants into the pool of practices on which selection operates. A niche's structure and its isolation from the introduction of new variants is partly constitutive of the degree to which scientific knowledge exhibits the directional trend we associate with scientific progress.

3.2.1 Isolation and Niche Construction in Kuhn

The importance of isolation for scientific progress was not lost on Kuhn, who had already noted in *Structure* that "Although it has become customary, and is surely proper, to deplore the widening gulf that separates the professional scientist from his colleagues in other fields, too little attention is paid to the essential relationship between that gulf and the mechanisms intrinsic to scientific advance" (Kuhn 1962b, 21).

There is something distinctively evolutionary about the association drawn here by Kuhn between the isolation of a group on the one hand and that group's ability to make progress on the other. In the same way that Kuhn asserts that the mechanisms that drive disciplinary progress can only operate under conditions in which the members of different disciplines are cordoned off from one another, the notion that a group's isolation is essential to its development as a single unit capable of directional evolutionary change is at the heart of the modern evolutionary theories of speciation and adaptation.

Yet it took another 30 years, in his presidential address to the Philosophy of Science Association in 1990, before Kuhn would draw an explicit connection between the role of isolation in fostering scientific progress and its role in the evolutionary process. Remarking on what he called the "unit of speciation," he wrote:

> In the biological case, it is a reproductively isolated population, a unit whose members collectively embody the gene pool, which ensures both the population's self-perpetuation and its continuing isolation. In the scientific case, the unit is a community of intercommunicating specialists, a unit whose members share a lexicon that provides the basis for both the conduct and the evolution of their research and which simultaneously, by barring full communication with those outside the group, maintains their isolation from practitioners of other specialities. (Kuhn 2000, 98)

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Kuhn's mature evolutionary view credits lexical diversity as "the isolating mechanism required for the development of knowledge" (Kuhn 2000, 98–99). Lexical diversity causes "breakdowns in communication," which end up limiting the "range of possible partners for fruitful intercourse." This "limited range of possible partners . . . is the essential precondition for what is known as progress in both biological development and the development of knowledge" (8). The ability of lexical diversity to delineate those with whom a scientists could and could not fruitfully engage was, in Kuhn's view, what gave rise to new specialized scientific disciplines, and it is these disciplines that host the process of epistemic refinement that Kuhn named *normal science*.

It was also in this venue that Kuhn introduced the concept of a biological niche to further elaborate his earlier assertions to the effect that scientists construct the world with which they are critically engaged. What Kuhn saw in the concept of a niche was a model for his earlier vision of a kind of feedback process in *Structure*. In this process, changes in the scientist's understanding of the contours of natural phenomena cause changes in how he approaches inquiry, which in turn results in further changes in his understanding of nature's contours, and so on. Just as niche-constructing organisms generate many of the factors that significantly affect their reproductive success, so too does the scientist's beliefs about what the world is like—"facts that the paradigm has shown to be particularly revealing of the nature of things"—determine to a significant degree how he designs his approach to inquiry (Kuhn 1962b, 21). For Kuhn, the components of an organism's niche and the way scientists conceptually represent nature play a "world-constitutive role" (103). The downstream results of the scientist's designated approach, as with the morphological evolution of organisms, is partly a product of which factors have been designated as relevant, and partly of the objective facts about how well certain structures—morphological or conceptual—perform in certain environmental conditions.

To summarize what I take to be unambiguous in Kuhn's final published thoughts on these matters: (1) isolation is an essential precondition of scientific progress; (2) lexical diversity creates "breakdowns in communication"; (3) communication breakdown isolates groups of scientists from one another; (4) a group's lexicon determines the representation of the world under which its members labor; (5) different representations propel inquiry in different directions; (6) inquiry so propelled is what gives rise to normal

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

science, a refinement process consisting of a feedback loop between scientists' representations of nature and their attempts to investigate it.

As I have already stated, for the same reasons that reproductive isolation is important for adaptation in biological populations, I believe that Kuhn was correct in assigning a critical causal role to limiting "the range of possible partners for fruitful intercourse" in the progressive development of scientific knowledge. One gets the sense that *The Structure of Scientific Revolutions* would have been a vastly different book had Kuhn noticed this connection at an earlier stage in his thinking. I also agree in part with components (2)–(6), with one essential difference—namely, that lexical diversity is, in my view, just one among many potential mechanisms through which groups of scientists become isolated from one another. I further deny that different specialized groups are characterized by different representations of the world, nor is this the most significant difference between them. By implication, I take the different lexical representations under which different groups sometimes labor to be just one—and by no means the most significant—of the factors promoting divergence between groups.

3.2.2 Isolation and Niche Construction beyond Kuhn

I want to begin outlining my more general picture of niche construction and isolation in science by picking up at the point where Kuhn's evolutionary considerations seem to have rested. Kuhn rightly emphasizes the dependence of a group's ability to make scientific progress on its isolation from other groups. But there are two fundamental mistakes in his approach that prevent him from a full appreciation of the affinity between the development of science and the evolutionary process. The first mistake is that he thinks that inability to communicate is the primary means by which groups of scientists become isolated from one another. The second mistake is that he thinks lexical diversity is the only way of generating communication breakdown.

Let us first look at Kuhn's esteem for the evolutionary significance of the inability to communicate. Kuhn, it seems, was inclined to model the isolation of scientists on a geographic theory of speciation that requires the presence of an impassable physical barrier to obstruct gene flow between populations. Like physical barriers, he saw the inability to communicate as necessary for isolating groups of scientists from one another because it makes the exchange of ideas all but impossible. Kuhn, I believe, was

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

committed to the notion that idea exchange (like gene exchange in the geographic theory) had to be impossible in order limit the number of partners available for fruitful interaction, because he (1) understood speciation to require the impossibility of gene flow, and (2) understood physical barriers to be required for making gene flow impossible. This view is expressed in the final pages of Kuhn's last published statement of his views:

> What permits the closer and closer match between a specialized practice and its world is much the same as what permits the closer and closer adaptation of a species to its biological niche. Like a practice and its world, a species and its niche are interdefined. And in both cases, also, that interdefinition appears to require isolation: the increasing inability of the residents of different niches to crossbreed, on the one hand, and the increasing difficulty of communication between the practitioners of different specialties, on the other. (Kuhn 1993, 337)

This view of how isolation fuels speciation is misconceived. It is true that some form of isolation has long been considered an important factor in facilitating adaptive evolution. Isolation was, for example, at the heart of Darwin's early, pre-*Origin* picture of adaptive speciation (Ospovat 1983, chap. 2). As Kuhn knew, the Modern Synthetic theory of speciation revolves around the geographic isolation of subgroups, owing to the central role given by that theory to the cessation of gene flow as part of the process by which subgroups diverge from one another. But recent scholarship has called the necessity of gene flow cessation into question, even if the role of cessation in facilitating divergence and in intensifying its pace and degree do not appear to be in doubt. In short, speciation does not depend on the impossibility of gene flow. It depends on genetic divergence, which can occur despite gene flow (Nosil 2008; Feder, Egan, and Nosil 2012). As a recent survey of the literature concludes, "the general conclusion emerging from . . . theoretical work is that . . . strong selection can overcome moderate gene flow to result in population divergence" (Nosil 2012, 146).

Similarly, specifically geographic forms of isolation like physical barriers are now widely recognized as a contingent, even if dominant, means by which subgroups become reproductively isolated from one another. Physical barriers are just one instance of the sort of thing that causes genes in a mating pool to cluster in certain ways. As mentioned above, differences in mating preferences, for example, have long been seen as capable of underwriting the kinds of genetic divergence characteristic of speciation, because they effectively segregate the mating pool into distinct subgroups. When

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

females strongly prefer not to mate with certain kinds of males, they don't need a physical barrier separating them from the males they find unattractive. They just avoid mating with those males.

The great irony in all of this is that one of the major achievements of *Structure* was how it drew attention to and outlined the many mechanisms by which groups of scientists often become isolated from one another—and most of these mechanisms have nothing to do with the inability to communicate. Rather, the isolation of scientists in Kuhn's early picture merely involves the adoption of a certain set of commitments—a *paradigm*—by a sufficiently large group of people. These commitments have normative consequences, in the sense that failing to subscribe to them places one outside the scientific community; in his words, they "define the legitimate problems and methods of a research field for succeeding generations of practitioners" (Kuhn 1962b, 10). The following passage is exemplary of what Kuhn saw as the kind of normative role performed by paradigms:

> One of the things a scientific community acquires with a paradigm is a criterion for choosing problems that, while the paradigm is taken for granted, can be assumed to have solutions. To a great extent these are the only problems that the community will admit as scientific or encourage its members to undertake. Other problems, including many that had previously been standard, are rejected as metaphysical, as the concern of another discipline, or sometimes just too problematic to be worth the time. A paradigm can, for that matter, even insulate the community from those socially important problems that are not reducible to puzzle form. . . . Such problems can be a distraction. . . . One of the reasons why normal science seems to progress so rapidly is that its practitioners concentrate on problems that only their own lack of ingenuity should keep them from solving. (Kuhn 1962b, 37)

Similar remarks are made regarding the normative status of certain metaphysical commitments, solution forms, and attitudes toward inquiry (Kuhn 1962b, 37–42). Like being part of a preference-governed mating pool, being part of the scientific community is not simply a matter of being in the right place at the right time. One must satisfy the community's norms for inclusion; failure to do so leads rapidly to one's exclusion from opportunities for "fruitful interaction." One needn't lack the ability to communicate in order to be excluded. All one has to do is run afoul of the prescriptive scope of paradigm commitments.

To the extent that Kuhn initially saw a deep connection between paradigm-driven science and the isolation of scientific communities, that

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

connection resided in the normative status that accrues to practices that have become sufficiently widespread. Whatever position one takes on the existence or nature of Kuhnian paradigms, there can be no doubt that scientific communities often exhibit strong preferences at any given time regarding what is a worthwhile research problem, which methods are reliable, how plausible a given theory is, and so on. These preferences are the barrier through which practices must pass in order to gain entry into the pool of partners available for fruitful interaction.

In this respect, the evolutionary significance of scientists' preferences regarding scientific practice strongly resembles that of mate preferences. Like mate preferences, the scientists' preferences regarding scientific practice isolate the population from the introduction of new practice variants—new problems, new theories, new methods, and so on. Whenever scientists' preferences are generated internally by members of the scientific community, they, like mate preferences, are apt to exhibit an arbitrariness with respect to the external environment that allows them to remain stable across a range of environmental circumstances. As Kuhn notes above, they "can . . . even insulate the community from those socially important problems that are not reducible to puzzle form." In this way, the preferences of community members can thus possess two of the key components involved in facilitating directional evolution: (1) the power to isolate the community from the introduction of new variants; and (2) the power to stabilize the pressures to which scientific practices are adapting.

Just as Kuhn's view of the isolating mechanisms in science was too narrow, so his conception of a scientific community's status as a constructed niche leaves out most of the ways in which scientists themselves can affect the reproductive success of practices. Earlier we saw how Kuhn used the niche concept strictly in reference to the fact that scientists conduct inquiry according to a certain representation of the phenomena; different groups of scientists occupy different niches to the extent that they represent phenomena in different ways. This is correct as far as it goes: a group's adherence to a particular perspective on nature is one of the factors that affect the reproductive success of practices in that group. Attempts to introduce practices that conflict with that perspective will sometimes be met with the sort of befuddlement and rejection that Kuhn describes in *Structure*'s chapter 6 on anomaly, and the propensity for members of the scientific community to react in this way can have an effect on the tendency for that practice

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

to become widespread. But there is no reason to think that conflict with a conceptual framework is the only way in which scientists themselves exercise control over a practice's reproductive success. Kuhn himself describes several. To take one example:

> I know of no case in the development of science which exhibits a loss of quantitative accuracy as a consequence of the transition from an earlier to a later theory. Nor can I imagine a debate between scientists in which, however hot the emotions, the search for greater numerical accuracy in a previously quantified field would be called "unscientific." Probably for the same reasons that make them particularly effective in creating scientific crises, the comparison of numerical predictions, where they have been available, has proved particularly successful in bringing scientific controversies to a close. Whatever the price in redefinitions of science, its methods, and its goals, scientists have shown themselves consistently unwilling to compromise the numerical success of their theories. (Kuhn 1961, 185)

Let us assume *arguendo* the much weaker claim that increases in quantitative accuracy often mark the transition from one scientific theory to another. If we further accept that this is no accident—that what explains this trend is that, in certain cases, members of the scientific community simply will not allow the ascendency of a less accurate theory—then this seems to be an unambiguous instance of scientists exerting control over the fate of theories. They have done so merely by demanding that new theories at least maintain the current theory's degree of accuracy. And as Kuhn states, the preference for numerical accuracy does not appear to be perspectival in the least; it can be exercised across a range of different representations of the phenomena.

The general point is that it is within the purview of a scientific community to decide by arbitrary fiat which sorts of practices will be considered for introduction into the pool of scientific practices and which will not. The exercise of normative constraints on proposed practices needn't solely be the consequence of a certain way of conceptualizing nature. The community can decide on whatever normative constraints they choose, and they need only to effectively enforce them in order to prevent the infiltration of things they don't like. Many modern scientific communities have chosen to insist on the preservation of established levels of numerical accuracy; they might have decided otherwise. Had they done so, Kuhn's generalization about increasing accuracy across the history of science would probably not have been true (if even it is true). Indeed, they might have decided on something far less intuitively appealing—say, that new theories

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

must have lower numerical accuracy than their predecessors. Or that new theories must be more whimsical than their predecessors. It doesn't matter. Whatever they choose will constrain the introduction of new practices in precisely the same way as our concepts—namely, it will oppose the introduction of new variants that conflict with it. In point of fact, arbitrarily designed rules for inclusion will probably be much more effective than practitioners' representations at resisting the introduction of incompatible new variants. After all, human groups are much more successful than biological populations at preventing the introduction of new variants into the mating pool (Laland, Odling-Smee, and Feldman 2000, 143). Whereas the ability of biological populations to resist gene flow often depends on factors quite beyond their control, the range of available practices can be policed effectively with minimal effort through conscious search and detection.

Scientific communities in fact employ an array of mechanisms for isolating the pool of scientific ideas from the deleterious effects that variation tends to have on the adaptive process. Some of these have been canvassed in chapter 2. The pedagogical practice of confining the content of undergraduate education to the use of a set of problem-solving strategies whose productive utility has been amply demonstrated functions as much as a tool for the suppression of novelty as it does as a way of preserving the favorable developments in theory and method from one generation of scientists to the next. The same is true at the graduate level, where burgeoning professionals tend to have little influence on what questions get asked and on how to answer them.

The efforts of scientific communities to prevent the introduction of new variation are no more clearly on display than in the peer review of research funding applications. The guidelines for how to decide whether a research proposal should receive funding include everything from general advice (such as that a proposal should build on existing work in a given field) to specific calls for research on very narrow problems. For example, a recent National Institutes of Health (NIH) Request for Applications reads:

> The purpose of this Funding Opportunity Announcement (FOA) is to support studies to evaluate candidate cellular therapies to treat radiation-induced injuries in appropriate in vivo models when administered 24 hours or more post-irradiation. The research should be aimed toward the eventual approval of candidate cellular therapies under the United States Food and Drug Administration (FDA) Animal Rule licensure pathway.[3]

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

The availability of funds devoted specifically to this problem will (it is hoped) increase the frequency of research projects aimed at addressing it. This is a deliberate attempt to influence what sorts of variants are available in a given generation, and there is reason to believe that it is effective. It is no accident that massive governmental support for the development of an atomic weapon led rapidly to the development of an atomic weapon.

Peer review of research publications is the second line of defense against the introduction of unwanted new variation. Though there are important differences in the sorts of considerations that determine success in the different contexts, publication review can be seen as a scientific community's last chance to prevent new practices from becoming available for reproductive opportunities. Here, as in research proposal review, members of a scientific community exercise their expert understanding of whether the research is important enough to warrant the community's attention and whether the evidence reported is sufficient to warrant the community's trust and acceptance. I would not deny that adherence to a particular conceptual framework may affect these decisions. Indeed, the NIH gives an explicitly Kuhnian affirmation of the high probability of such an event: "be cautious about seeming too innovative. Not only is innovation just one of five review criteria, but a paradigm shift can be as dangerous as an earthquake! A reviewer may take a challenge to the status quo as a challenge to his or her world view."[4] Rather, what I deny—and what I think Kuhn would have also denied, had he been asked—is that new practice variants cannot be excluded due to commitments far more modest than those involved in resisting conceptual change.

These considerations indicate ways in which a scientific community can come to constitute a kind of cultural niche, a niche constructed out the prevailing attitudes of community members and isolated from exogenous influences—even from other scientific communities. By policing the avenues through which new practices get introduced into the pool of practices available for fruitful interaction, scientists' preferences function as a mechanism for isolating scientific practice from deleterious new variation—variation that would otherwise undermine the directional gains made by previous generations. The decision about which preferences are active in restricting the introduction of new practices is entirely within the hands of the scientific community. Numerical accuracy cannot simply shove its way into the pool of practices; its introduction and maintenance there depends

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

solely on the community's attitude toward it. In this way, community-level preferences about permissible approaches to inquiry strongly affect the reproductive success of practice variants. As Kuhn (1962b, 92) put it, "there is no standard higher than the assent of the relevant community."

I have been surveying some of the ways in which scientific communities manufacture the determinants of reproductive success of scientific practices. Notice that this is not to say that communities manufacture the determinants of all types of "success." The mind-independent world contributes to whether our theories are, for example, predictively successful. But it is scientific communities that must make the decision to care about predictive success, and to use it as a criterion by which they measure the desirability of practices. And this is what matters to whether a practice is increasing or decreasing in frequency.

We are now in a position to say that, like the niche-constructing biological populations that manufacture the factors that affect reproductive success, scientific communities facilitate the directional refinement of scientific practice through their ability to stabilize selective pressures across a wide range of environmental circumstances. And we can understand the notion of "environment" here to include human social as well as physical and biological circumstances. To echo Kuhn's observation above, this includes environments in which a problem achieves a high degree of social importance. Whatever might be going on outside a given scientific community, the decisions regarding which practices to admit into, retain within, and eject from the pool of variants lies at all times within the discretion of the scientific community. Preferences regarding the composition of the variant pool may be exercised with more or less efficacy, and, as Kuhn emphasizes, more or less cognitively. But, for communities that stabilize their practice preferences across different environmental circumstances by isolating themselves from the broader society's estimation of a problem's importance, the process of directional refinement is largely within their control. No one familiar with the menu of problems in contemporary mathematics or particle physics will be surprised to hear that a problem's social importance normally has little influence on whether it is important to practitioners. As predicted by the evolutionary considerations outlined above, history suggests that disciplines that isolate themselves effectively tend to oversee more directional refinement in practice than disciplines

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

that are responsive to external pressures. Free from the influx of unwanted intellectual contamination, a well-isolated scientific community can effectively push scientific practice toward an internally defined optimal state. Ultimately, it is an empirical matter whether any particular community is able to isolate itself and stabilize the determinants of practice propagation in a way that facilitates directional refinement. And, of course, the notion that the cultural norms of scientific communities affect the content of practice is not new. What I hope to have done here is to have drawn out more explicitly the connection between (1) these norms and the isolation of practice, and (2) the isolation of practice and the propensity toward directional refinement. The point of this section has been to establish the theoretical point that scientific communities that isolate themselves and stabilize pressures in this way create the sorts of conditions that make directional refinement possible.

3.3 A Niche Constructed out of Epistemic Pressures

In this section I make the case for when we should understand a scientific community's isolation mechanisms to be epistemic phenomena—that is, phenomena whose existence in science is explained by their distinctive contribution to improving our knowledge. Contrary to what is often assumed in science studies literature, the mere fact that scientific communities have an identifiable culture does not by itself warrant the assumption that the culture of scientific communities is indistinguishable from that of other communities with respect to its ability to move to epistemically superior states. Not all cultural arrangements are equal in their capacities to solve difficult problems, nor are all scientific communities. Some communities happen to be blessed with an unusually high concentration of cultural mechanisms that promote problem-solving.

3.3.1 The Admissions Process for Problem-Solving Approaches

Certain scientific communities have in place an array of mechanisms for isolating the pool of scientific ideas from the deleterious effects that variation tends to have on the adaptive process. Some of these have been canvassed in chapter 2, where particular focus was given to the suppression of novelty. The key to understanding how the suppression of novel problem-solving approaches facilitates adaptation within the scientific community lies in the way the shared commitment limits the means available to practitioners

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

for solving particular problems. Through their professional training, practitioners develop an intuition for seeing a particular problem as a species of some genus or other, some general class of problem for which a general problem-solving approach has been established. Given a particular case, for example, an experienced practitioner will see that it is best to "treat it as a [name of genus] problem." And, in general, the wide adoption of new problem-solving strategies is stringently resisted unless it possesses clear virtues over available approaches—virtues such as computational ease, generality (which reduces the number of strategies needed for a given specialty), or success in dealing with hitherto insoluble kinds of problems. This is the premium that novel approaches must pay in order for admission into the scientific community's toolkit. Notice that, although novelty is prized at administrative levels, it is strongly opposed by scientists in practice. What value is novelty per se to the scientist? Her training has already equipped her with an intuitive sense of the resources to which she can appeal for tackling certain kinds of problems. What need has she for a new resource that offers nothing beyond its novelty? As Kuhn observed near the end of *Structure*, "Novelty for its own sake is not a desideratum in the sciences as it is in so many other creative fields" (Kuhn 1962b, 169).

Why does this resistance exist in certain fields, and how does it facilitate adaptation? The resistance exists because it is preferable to treat every problem as the sort of problem one already knowns how to solve quickly, mechanically, reliably, and with negligible cognitive effort. Anyone with experience doing problem sets in mathematics or the quantitative sciences is familiar with the experience of finding the first few problems absolutely confounding, but quickly becoming adept at solving them; what initially takes several minutes is soon cut down to a fraction of that time, often with a higher success rate. As seemingly new kinds of problems are encountered, one looks for ways of conceiving of them as a species of the same genus of problems that one is already adept at solving.[5]

There is an interesting parallel to this process that occurs at the generational level. Much of cross-generational methodological stability in the sciences has to do with the training involved in the use of research technology, a category that includes not only scientific instruments but also proof styles in mathematics, argument styles in philosophy, ethnographic styles, historiographic methods, and much of what goes under the "research methods" headings in graduate programs. As is well known

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

by anyone trained to do research, professional mastery over a particular problem-solving approach is not something that can be achieved by reading an instruction manual. Rather, the sort of expertise that characterizes a practitioner of a given research specialty evolves by observing and assisting in the competent execution of an approach by previously trained experts. No research publication could possibly transmit the volume of knowledge required to successfully employ a research method, even one that is well known and widely employed. This is because so much of research involves the cultivation of an instinct for pattern recognition, an instinct that regularly defies propositional explication.

The methodological stability of a scientific community is thus partly a consequence of the nature of expertise. If one can gain mastery over a technique only by apprenticing with a master, then cross-generational methodological continuity is essential to generate the grade of expertise that we know to be required for solving research problems. Of course, like the list of research problems, it is not as if the corpus of methods will not evolve over time. But as I mentioned earlier, the introduction of new methods is generally resisted by practitioners unless there is a significant problem-solving payoff. And even when new methods are incorporated, as with any technology, it takes several generations of trained experts to really get things working in a well-oiled manner. The research generated by a discipline's initial foray into the new technology looks primitive when measured against that generated by students just one or two generations later, who "grew up" on the technology that their teachers frustratedly fumbled through. Being able to wield a problem-solving strategy with the degree of sophistication required to crack problems of respectable depth is a skill that scientific communities take several generations to develop. Permitting, let alone expecting, young researchers to innovate for innovation's sake would deprive the scientific community of the second- (and third-, and fourth-, etc.) generation inheritance of research methods that have reached a stage of development sufficient for handling challenging problems.

There are thus a variety of good epistemic reasons that scientific populations might aim to suppress superfluous novelty, each of which helps to promote problem solving. If we understand a scientific community's tendency to resist new problem-solving approaches as an attempt to maintain and refine its own cognitive efficiency, the tendency's ability to generate progress makes total sense.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

3.3.2 The Admissions Process for Scientific Problems

Let us now look focus on the filtration process for scientific problems. New problems are introduced into the community when they are deemed *significant*. Significance is one of the more-opaque quantities that plays a determinative role in scientific reasoning. Understanding its content is an important philosophical project, but it is not one that can be productively executed in the space of this chapter (I address it in chapter 5). As a form of triage, we can examine major scientific institutions' best effort to articulate what it means for research to be significant, which is as accurate as any account we currently possess is likely to be.

At the center of these institutions' attempts to gauge significance is the examination of the strength, proximity, and number of connections to *ongoing scientific research* (Haufe ms-b). This is directly in line with what we would expect to see in isolation mechanisms if they were designed to restrict a scientific community's problem-solving priorities to the specific set of problems whose significance has already been established by the community itself. The simplest means by which to restrict problem-solving priorities in this way is to prevent any new problem whatsoever from entering into scientific practice. For obvious reasons, this would be suboptimal. Restricting scientific research to problems that are either (1) old, significant ones or (2) new problems that are closely related to the old, significant ones allows the scientific community to introduce new problems by using their relatedness to old, significant ones as a filter for distinguishing potentially significant new problems from new problems that we have little reason to think will be significant. In any event, significance is conceived of strictly in terms of its relation to the scientific community's previously established commitments.

This approach to judging the significance of scientific problems is at odds with what we would expect isolating mechanisms to look like if significance were primarily a function of personal or broader sociocultural priorities. Rather than an exclusive focus on proximity to current research problems, a scientific community whose adaptive directionality was significantly affected by concerns external to knowledge would be expected to conceive of scientific significance in more general terms. To be sure, funding agencies generally do provide space for external concerns to be taken into account (e.g., "Broader Impact"), but they explicitly divorce these concerns from those that are genuinely relevant to scientific inquiry (Haufe ms-b). This idea bears emphasizing. The largest and most influential public

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

mechanisms for scientific research funding in the world—the NIH, the National Science Foundation, and the Research Councils UK—all make a distinction between factors that make research significant for science and factors that make research significant for society.

3.3.3 The Lack of Covariation between Science and Society

If the direction of scientific inquiry is strongly conditioned by the social context of inquiry, at the very least we would expect core aspects of inquiry to vary significantly with social context; scientists in similar social contexts should exhibit more agreement than scientists in disparate social contexts. Occasionally, we find that this connection does not strictly hold, which could indicate the presence of isolating mechanisms designed to effectively filter out factors that are external to the production of knowledge.

In this regard, there are four kinds of phenomena in particular that might suggest an epistemic role for some isolating mechanisms in certain scientific communities. The first three concern cross-cultural convergence with respect to (1) which scientific problems are significant, (2) which methods of inquiry are best, and (3) which empirical commitments are most well established. As numerous authors have shown, not only convergence itself, but the rate of convergence with respect to these three dimensions in science is most strongly affected by the rate of communication among scientists.[6] For example, as Bettencourt et al. (2006) showed, the rate at which scientists from the United States, Japan, and the USSR adopted the use of Feynman diagrams is virtually identical once some scientists from a given country had been introduced to them. The overwhelming effect of scientific communication on convergence is the expected result if the isolated entity is a community of scientists whose intellectual objectives are relatively immune to social context. This echoes Gerald Holton's observation that:

> Despite the variety, despite the specialization that makes it difficult to follow what is being done in the laboratory next door, despite the important differences between basic and applied, large and small, or experimental and theoretical physics, its practitioners still clearly conceive of themselves as doing in different ways work in one identifiable field. There are no large cleavages and disputes between sizable factions representing fundamentally different styles. (Holton 1962, 96)

Closely related to these is the phenomenon of "multiple discovery," whereby two individuals independently develop the same theory, uncover the same effect, prove the same theorem, invent the same technology, and

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

so on. Again, this is the expected result if the direction of inquiry is overwhelmingly a consequence of adaptation to research problems themselves, just as it is in evolutionary biology. But it would be a rather surprising occurrence on the assumption that influences of social context significantly affect inquiry's trajectory. For instance, the independent discovery of fast Fourier transformation by Gauss in the late eighteenth century and then again by Cooley and Tukey in 1965 is much easier to understand if we imagine the adapting entity to be one characterized by the pressure to reduce the computational difficulty involved in the analysis of a function, rather than pressure to accommodate a set of social objectives that is common across eighteenth-century Germany and the twentieth-century United States (Heideman, Johnson, and Burrus 1985). Were the latter interpretation correct, we would expect most of these multiple discoveries to occur within similar social contexts, which is not the case.

This by no means suggests that social factors per se are irrelevant to understanding the directionality of science. As I argued in previous chapters, adaptive problem solving—be it in the physical sciences or in biology—requires persistence and isolation of the problem-solving entity. The dominant factors responsible for the persistence and isolation of the scientific community are the various sociocultural institutions involved in education, professionalization, research problem selection, and problem-solving method selection. It is no more reasonable to think that scientific communities could evolve adaptively without these socially embedded enforcement mechanisms than it is to think that biological populations could evolve adaptively even if offspring tended not to resemble parents.

3.4 The Efficacy of Epistemic Pressures

Let us now use the points established in sections 3.2 and 3.3 to determine the conditions under which epistemic selective factors matter more to the development of scientific practice than do other evolutionary factors, such as drift and migration. The degree to which scientific practice approximates a paradigmatically Darwinian population depends on the degree to which the propagation of practices is determined by how well designed they are for their environment. And the degree to which the propagation of practices is determined by how well designed they are for their environment, in

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

turn, depends on the degree to which the selection process is shielded from the disruptive effects of other evolutionary processes.

I argue that a scientific community's status as a constructed niche contributes strongly to its ability to dampen the effects of forces that threaten the directional efficacy of natural selection. One of the reasons to think that this is the case has to do with how niche construction affects the magnitude of selection pressures. Another reason to think this is the case is simply the track record of solved problems to which many scientific communities can claim credit. I discuss each of these reasons in turn.

3.4.1 Niche Construction and the Magnitude of Selection Pressures

We can think of the alleles in an offspring generation as a sample of the alleles present in the parent generation. If the process by which that sample is generated is unbiased, and if each generation contains the same number of individuals, we expect the frequency distribution of alleles in the offspring generation to mirror that of the parent generation. Differences in adaptedness introduce a bias into the sampling process that makes better-adapted types more likely to appear in future generations. This sampling bias—what we refer to colloquially as *natural selection*—can affect the frequency distribution of types in a population according to their relative adaptedness. But natural selection is not the only means by which frequency distributions can be affected; there are factors other than relative adaptedness that contribute to a type's representation in a particular generation.

One of those factors is sampling error. We refer to sampling error in biological populations as "drift," and, like statistical sampling generally, certain factors affect the probability of its occurrence. When dealing specifically with evolutionary drift, what we are most interested in is how good of a guide an individual's adaptedness is to his representation in future generations. In the absence of sampling error, an individual's reproductive success is an excellent predictor by how well adapted he is to his environment. But the predictive value of adaptedness tends to decrease when the magnitude of selection is weak. This is because, under weak selection, the environment does not harshly discriminate among degrees of adaptedness. Under these conditions, nature does not exhibit a strong bias in favor of well-adapted types, and the sampling process becomes relatively indiscriminate. The probability of sampling error increases in such circumstances:

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

the degree of expected reproductive success conferred on an individual by virtue of his adaptedness (*fitness*) ceases to be a good predictor of his actual reproductive success.

As one might expect, things are precisely the reverse under strong selection. The more strongly nature discriminates on the basis of adaptedness, the better a predictor adaptedness is of actual reproductive success. The significant implication for our purposes is that increasing the strength of selection is one way of reducing the probability of sampling error. Thus, the strength of selection has something to tell us about how significant of a contributor natural selection is to the frequency distribution of types in a population—strong selection suggests effective selection, because strong selection reduces the probability of sampling error, and reducing the probability of sampling error increases the efficacy of natural selection.

Understanding scientific communities as niche constructors becomes relevant here because of the amplifying effect that acts of niche construction have on the magnitude of selection pressures. We can think of an instance of niche construction as a type of acute environmental perturbation. Much of the time, environments change at a rate that is slow enough to allow populations to evolve adaptively over several generations. But once in a while an environment may undergo some form of acute perturbation, such as a virulent illness, an invasive species, or abnormal weather conditions. Because of how rapidly the environment changes, the selective effect of an acute perturbation tends to be much stronger than it would be had it been introduced more gradually. The great paleobiologist Dave Raup referred to this kind of event as a "first strike," and described it as one of the necessary ingredients for driving a population to extinction (Raup 1991).

Acts of niche construction are a type of acute environmental perturbation in the sense that they are capable of significantly altering the selective environment in a relatively short time frame. When niche constructors establish some factor as relevant to survival and reproductive success, the sudden, qualitative shift in what constitutes a sufficient degree of adaptedness generates at the outset intense selective pressure on members of a population. As with other forms of acute environmental perturbation, this initial phase is then followed by a relative relaxation of selective pressure exerted by the perturbative element. Relaxation occurs because those individuals who were most sensitive to the selective effects of the perturbation

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

were culled from the population in the initial phase, leaving alive only the less sensitive individuals.[7]

This, finally, is where support for the Intrinsicality Thesis starts to come into view. We began this chapter by asking to what degree differences in the rates at which scientific practices copy themselves can be explained by differences in their adaptedness to a given environment. And we've just established that acts of niche construction tend to increase the magnitude of selection in a very short time. Now, it is generally understood that the degree to which differences in adaptedness explain differences in reproductive success is determined by the product of population size (N) and the magnitude of selective pressure (s) (Roughgarden 1979, 74–79). Thus, niche construction must increase the degree to which differences in adaptedness contribute to reproductive success, because it increases the magnitude of selective pressure. Now, if and when scientific practices occupy a constructed niche, differences in the value of different scientific practices tend to affect the development of scientific practice in predictable ways. And, if the practice niche is constructed out of epistemic pressures, then we can at least say the following:

> Differences in the reproductive success of scientific practices are more likely to be caused by differences in how well suited those practices are for promoting the growth of scientific knowledge than if scientific practices did not occupy a constructed niche.

To the extent that a scientific community qualifies as a cultural niche, this result is fairly secure. But it is not overly impressive. We can, however, strengthen it by noting conditions under which the selective pressure exerted by scientific communities is strong not just by comparison with nonconstructed niches, but also in an absolute sense. Are there any such conditions?

In section 3.2.2 we looked at the research funding review process to get a sense of the nature of certain communities' preferences regarding which research problems to pursue and how to pursue them. Focusing just on this layer of the selection process suggested that the preferences were overwhelmingly aimed at the pursuit of epistemic goals; they are the kinds of preferences one would expect to see exercised if the goal of the review process were solely to promote epistemic ends. We can now return to the research funding review process to learn more about how strongly these preferences are exercised, because we actually have quantitative data on

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

how influential different criteria are for assessing the overall quality of proposed research (Eblen et al. 2016; Lindner, Vancea, and Chen 2016).

The first thing these data tell us is that most of the difference between proposal scores is explained by (1) differences in how highly reviewers esteem the research methodology and (2) differences in how highly reviewers esteem the research questions. The other criteria—environment, investigator, and innovation—have relatively little influence. This distribution of weightiness is precisely what one would expect if the process were designed to promote epistemic ends. We associate improvements in scientific knowledge primarily with whether we're asking the right questions and whether we're answering them in the right way. And these are the considerations that dominate the review process.

The second lesson to be learned from the data is that relative to other factors potentially affecting the overall quality assessment, the magnitude of the pressure exerted by the "Approach" and "Significance" considerations is extraordinary, with "Approach" being almost twice as strong of a predictor as "Significance" of overall assessment. The next strongest predictor is an application's "Innovation" score, which is less than half as influential as "Significance." Both the predictive magnitude and rank-ordering of these criteria suggest that (1) the more directly a criterion measures epistemic properties of the proposal, the more influential that criterion will be; and (2) the most epistemic of the measures are strong determiners of an application's success.

Perhaps equally interesting is the fact, noted by Lindner, Vancea, and Chen (2016, 246), that only expert peer reviewers seem to pay attention to the "Approach" score when generating their overall assessment. Recalling our distinction made in chapter 1 between internal and external pressures, I defined a pressure as internal to science *if and only if adaptation to it is explained by achievements that are recognizably important to most practitioners, but not to most nonpractitioners*. Now, since "Approach" is the closest thing to a purely epistemic feature of the proposed research, and it is something that—according to the Lindner, Vancea, and Chen study—only expert peer reviewers care about, then there is some relatively clear empirical warrant, spanning a wide range of scientific communities, for the thesis that adaptation to pressures internal to science is characteristically achievable only through improvements to our knowledge.

Admittedly, it is important not to overestimate the general significance of the research proposal review process for understanding how scientific

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

inquiry develops. But it is undeniably a key element of the process by which modern scientific knowledge grows. And, as a context in which one is able to directly examine the sorts of considerations that scientists think are important for distinguishing between more and less promising ideas, it is actually quite useful. While a study of the proposal review process cannot provide details on the magnitudes of all the different selective pressures, it does provide a sound empirical foundation for assigning magnitudes to some of the most significant selective pressures involved in shaping the composition of scientific practice across a broad range of disciplines.

3.4.2 Problems Solved

A final consideration that weighs in favor of the notion that epistemic selection pressures are effective in some scientific communities is the fact that, somehow or other, the problems that these communities set for themselves do tend to get solved. In biological contexts, we expect populations to evolve an adaptation to deal with a certain ecological factor only if that factor is a significant determinant of organism's reproductive success; if dealing with that factor is insignificant for reproductive success, there is no reason to expect populations to evolve a response to it. Scientists routinely formulate problems and develop theories that eventually solve those problems. If the determinants of success within these communities were independent of solving those communities' chosen problems, we should expect problem solving to be far less rapid and routine than it often is.

It might be objected that scientific communities are ultimately driven by pressure for prestige, and scientific problems get solved because that is the most effective way of gaining prestige. Thus, the driving factors behind the adaptive evolution of the scientific community are only indirectly related to problem solving. What ultimately explains the directionality of scientific inquiry is selection for prestige.

There are several mutually reinforcing ways of responding to this objection. First, if in fact there is pressure on scientists to solve problems because solving problems is the only way of gaining prestige, then we would still expect scientific inquiry to be organized in a way that maximized problem-solving effectiveness.[8] Thus, the fact that the distinctive culture of scientific communities is designed to maximize the prestige of individuals cannot undermine the special epistemic status of science if the only way to

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

maximize prestige is for individuals to move the community to epistemically superior states.

Second, the appeal to prestige (or similar social commodities) as the chief determinant of the direction of scientific inquiry does not actually add anything in terms of explanatory power when the chief determinant of prestige is solving problems. In other words, the reference to prestige is often a superfluous add-on that leaves completely unaltered our expectations for the trajectory of scientific inquiry, were we to merely assume that the adaptive significance of problem solving is unmediated by its connection to prestige. This example from the careful and considered work of sociologist Michael Mulkay is illustrative of the explanatory impotence of prestige: "They [scientists] tend to move into areas which appear to offer special opportunities for productive research, for the utilization of their particular skills and, *consequently, for career advancement*" (Mulkay 1975, 5; emphasis added). If we adopt the simplest, most conservative working hypotheses—that scientists "tend to move into areas which appear to offer special opportunities for productive research" because their goal is to generate productive research, or that they "tend to move into areas which appear to offer special opportunities . . . for the utilization of their particular skills" because their goal is to use their particular skills—then we have a complete, coherent, and intuitive explanation for the empirical facts as cited. No additional insight is gained by invoking prospects for career advancement.

Third (and last), the explanatory appeal to prestige treats it as uncontroversial that human action is explained by the pursuit of prestige. Although that principle may indeed be uncontroversial, the appeal to prestige as a motivating factor in science does not explain why problem solving itself is so important for prestige in many scientific communities. Why problem solving? Why not something else, like good looks, a winning personality, being rich, and so on? These factors are certainly sufficient for prestige in lots of other social contexts. Furthermore, why *only* problem solving? Why not problem solving plus something else? True, that something else has to be associated with prestige in science. But unlike domains in which problem solving is not linked to prestige, science has a widely acknowledged reputation for making epistemic progress. Either (1) it's a remarkable coincidence that the only domain in which problem solving is the routinely recognized as the dominant means by which prestige is achieved happens to also be the domain most closely associated with the notion of epistemic

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

progress; or (2) problem solving is often an effective means by which prestige is achieved in science, because problem solving is causally responsible for the directional phenomena picked out by our notion of epistemic progress, and epistemic progress is what really matters to many scientific communities. My money is on explanation (2).

3.5 Recap and Conclusion

We started out wanting to know what it would look like if intrinsic features of scientific practices were the most likely explanation for differences in the rates at which those practices copied themselves. What have we achieved at this point? First, we've come to understand the general importance of stability and isolation for facilitating directional evolution, and of the particular talent of niche constructors for promoting stability and isolation. We then looked at ways to think of scientific practices as occupying a niche composed of epistemic selection pressures (i.e., selection pressures that take the form of practitioners' preferences for certain kinds of practices over others). Finally, we considered situations in which we seem to have good reasons for believing that these selection pressures matter a lot to the rate at which practices copy themselves.

Most of what we've achieved in this discussion has been the development of a deeper understanding of what scientific practice would have to be like if it were the case that differences in intrinsic features of scientific practices cause most of the differences in reproductive success. And it's important to emphasize that this chapter, like previous and subsequent chapters in this book, is first and foremost an attempt to articulate what scientific practice would have to be like if it were a paradigmatically Darwinian population; in no way do I think that scientific practice in every discipline satisfies the constraints outlined above. The point is rather to describe an idealized state with the understanding that, the closer scientific practice is to that idealized state, the greater the degree to which it is able to partake of the directional power of natural selection.

Secondary to this point has been the goal of providing reasons to think that scientific communities satisfy these constraints to some degree in certain cases—that is, that there are broad similarities across many disciplines with respect to how closely they approximate idealized epistemic niches. This goal is difficult to attain without lots of different kinds of empirical

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

data, which, I'm sorry to say, is simply not available. But the two kinds of data we do have are promising. The quantitative data on research proposal review from the NIH sends a very strong signal in direct support of the generalized dominance of epistemic pressures across many different scientific disciplines. And the track record of problem-solving in many disciplines is difficult to reconcile with the notion that intrinsic features of scientific practices routinely fail to exert a strong influence over the rate at which those practices are copied.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

4 Directionality and the Refinement of Scientific Practice

4.1 Practices of Extreme Perfection

In one of the most celebrated passages in the *Origin*, Darwin confronts what he takes to be a critical challenge to his claim that adaptations arise through the gradual refinement of structures. The challenge was to show how the class of phenomena he describes as "organs of extreme perfection and complication" could possibly have arisen through the accumulation of "numerous, successive, slight modifications" (Darwin 1859, 189). It is in this context that the terrifying explanatory power of the principle of descent with modification through natural selection comes into sharp relief. In case after case, Darwin shows how we might come to see such wonders of nature as camera eyes, electric organs, and sterile insect castes as anything but miraculous; as rather little more than the expected result of a process of selective retention of favorable variations, carried out over "millions on millions of years; and during each year on millions of individuals of many kinds" (189).[1] In other words, we need appeal neither to the special creation of such organs ex nihilo nor to the equally improbable notion that the organ appeared fully formed as a natural variant. We need only "suppose each new state of the instrument to be multiplied by the millions; and each to be preserved till a better be produced, and then the old ones to be destroyed . . . and natural selection will pick out with unerring skill each improvement" (189).

Darwin admitted that, although in most cases we are ignorant of the ancestral forms of these organs, the several cases where the ancestral form is known contain a valuable general lesson—namely, that the adaptive function of the ancestral state need not match that of its present state. This lesson is important because it gives us insight into the solution to "an even

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

more serious difficulty," that is, "cases of two very distinct species furnished with apparently the same anomalous organ" (1859, 193). How is it that species as distinct as humans and cephalopods, separated by roughly a billion years of evolutionary history (having diverged some 500 million years ago), might nevertheless come to possess eyes that are highly complex, highly perfected, and work in basically the same way? Darwin answers:

> I am inclined to believe that in nearly the same way as two men have sometimes independently hit on the very same invention, so natural selection, working for the good of each being and taking advantage of analogous variations, has sometimes modified in very nearly the same manner two parts in two organic beings, which owe but little of their structure in common to inheritance from the same ancestor. (Darwin 1859, 194)

Through two very different sequences of "numerous, successive, slight modifications" preserved by natural selection, humans and cephalopods have managed to find their way to the same obscure point in morphospace, "in nearly the same way as two men have sometimes" found their way to the same obscure point in idea space. This is what contemporary evolutionary biologists call *convergent evolution*. And, as Darwin suggests, the same phenomenon is a routine occurrence in the context of human discovery and invention.

Convergent evolution is a special case of the general tendency of lineages to exhibit directionality under evolutionary pressure by natural selection. When two distinct lineages are under the same kind of pressure from their environment, the results of their independent attempts to respond to that pressure occasionally converge, because it sometimes happens that there is a particular response to that pressure that is (1) better than other responses and (2) achievable by both lineages through some finite series of steps. Under these conditions, the starting point, sequence ordering, and the specific nature of the changes can be rendered irrelevant to where the lineage ends up in morphospace. And when these same conditions are satisfied for two distinct lineages, they can end up much closer together than we might expect, given their vastly different evolutionary histories.

Convergent evolution is a striking illustration of natural selection's ability to push a population toward a particular point in morphospace, in that it indicates a certain *robustness* to a lineage's eventual occupation of that state. After all, if lineages as distinct as cephalopods and humans have come to occupy the same region in morphospace through natural selection, then

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

we have strong reasons to think that humans would have ended up there had our own evolutionary history been very different. The natural history of cephalopod eyes can be thought of as an alternative history of the camera eye—one in which it developed through a dramatically different pathway than its actual history.

In this chapter I describe in detail a model for understanding how central aspects of the development of scientific knowledge might be, as Darwin seems to have thought, governed by the same evolutionary process that leads a lineage to a nearly unavoidable point in morphospace. I argue that this process is responsible for the often inexorable transformation of our relatively vague and general ideas about nature and its investigation into we might call, after Darwin, practices of extreme perfection. I emphasize that evolution does not always work in this way—whether it be in biological or scientific contexts—and that history *does* matter for many things. But there are some evolutionary phenomena for which history does not matter, and those phenomena are of supreme importance with respect to our understanding of how scientific knowledge develops over time in a number of disciplines.

My argument builds on where Darwin left us, that is, with the empirical fact that two or more people often hit on the same idea or invention. I argue that this frequent occurrence in science is due to the fact that, as an isolated population subject to modification through natural selection, scientific practice often evolves directionaly toward an optimal state. These optimal states are a kind of *attractor*, a well-defined state toward which a system develops on account of certain general features of its governing dynamics.

I begin by looking at details of the phenomenon of multiple discovery in science, followed by an attempt to situate that phenomenon within the broader class of systems that reliably produce certain kinds of well-defined states. I then turn specifically to the subset of such systems that are governed by natural selection. Here I appeal to the notion of an adaptive landscape as a way of representing how populations ebb ever closer to an optimum state, looking at the various kinds of conditions that make that possible. In particular, we will want to look more closely at the significance of Godfrey-Smith's "Continuity" condition. Following that, I elaborate on what sorts of conditions scientific practice would need to satisfy, were it to

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

be maximally susceptible to adaptive optimization in the manner of paradigmatically Darwinian populations.

The adaptive landscape and its constituent concepts will give us the opportunity to see how different stages in the development of scientific practice can be characterized by differences in the extent to which scientific practice exhibits the criteria for what counts as a paradigmatically Darwinian population. As emphasized by Kuhn (1962b), the development of scientific knowledge is episodic, and the rate and nature of scientific progress differ dramatically from stage to stage. I try to show how these stages can be mapped onto an adaptive landscape in an explanatory way. Most importantly, we can use the general principles governing an "adaptive walk" to explain why different stages of the development of scientific knowledge often exhibit these characteristic differences.

As in previous chapters, this last step will involve viewing certain elements of the nature of scientific inquiry as analogues of the evolutionary process in biological contexts. But in this chapter we confront an additional complication, one that stems from the normative dimension inherent in our discussion of the directional refinement of scientific practice. In chapter 1 I argued for replacing the notion of scientific progress with one of directionality. This avoids some of the uncomfortable normative implications of the concept of progress, but not all of them. For we are still left with the claim that, through a process of directional refinement, scientific practice is in some sense "optimized"—and optimization is, at least in part, a normative notion. I will argue for a particular way of understanding the optimization of scientific practice that I think successfully avoids problematic implications while still retaining the ability to validate certain states of scientific practice as better than others.

4.2 Convergence in Practice Space

We can imagine a state space consisting of all the scientific practices one could possibly adopt—all possible scientific theories, questions about nature, methods of investigation, canons of reason, and so on. Each of these categories represents a dimension of practice space. Different points in the space represent different possible combinations of practices in various states. Call these different possible combinations *practice clusters*. Clearly, practice space is unfathomably large and contains continuum-many practice clusters.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Despite how roomy practice space is, the occupied points—the practice clusters that are actually employed by members of a given scientific community—are not evenly distributed. In fact, they are very, very densely packed together. Nor is there an even distribution of these groups of occupied points. They too tend to clump together in the same neighborhood in practice space. For some sciences, like physics, almost all of practice space is entirely empty.

High-fidelity copying alone cannot account for sustained clumpiness, only for initially clumpy conditions. Rather, the clumpiness of practice space suggests that movement within it is severely restricted—that is, that there are strong constraints on a practitioner's ability to adopt practices that are sufficiently different from the ones with which he began. For the moment, let us ignore what these constraints might be. Because practitioners start very close together, and because their movement is severely restricted, we would expect the movement of multiple practitioners to the same point in practice space to be a fairly routine occurrence. And indeed it is.

Remarking on how frequently these instances of what he called "multiples" occurred, and how frequently this fact had been pointed out, the sociologist of science Robert K. Merton went so far as to promote the radical thesis that all scientific discoveries were "in principle multiples." By this, Merton meant that even discoveries that are in fact known to have been made by only one person are "potential multiples that happened to emerge as singletons." On his view, it is "singletons, rather than multiples, [that] are the exception requiring distinctive explanation" (Merton 1961, 478). In other words, scientific inquiry, by its very nature, drives independent practitioners to the same point in practice space. Convergence is the default expectation, and deviations from that default expectation (like singletons) demand consideration of causal factors that are not endemic to scientific inquiry itself.

One of the pieces of evidence Merton marshaled in support of his radical thesis was the fact that we frequently find that discoveries we thought were singletons were actually rediscoveries of results derived previously by other researchers. Merton cites Henry Cavendish and Carl Friedrich Gauss as "instances of a larger class" of researchers who often did not publish their results, and whose privately recorded discoveries were routinely derived independently by subsequent generations. Closely related to this are the common occurrences whereby (1) a scientist learns shortly before going to

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

print that he is not the first person to have made the discovery or (2) he learns that he has been scooped and consequently decides not to submit his results. Merton took this thesis and several other similar phenomena to be commonplace across the history of modern science and symptomatic of the convergent nature of scientific inquiry.

The other class of evidence to which Merton appealed was the various effects that knowledge of the convergent nature of science ended up having on scientific inquiry itself. Merton was drawn to the study of multiples by his interest in the sociology of priority disputes in science, which, as Merton perceptively pointed out, can only happen if multiples occur. One of the many interesting observations he made about priority disputes is that they began to decline precipitously in the nineteenth century—even though multiples did not. According to Merton, disputes over priority of discovery began to die down once scientists started to realize that multiples are a routine part of scientific inquiry—the "expected result," to use Merton's term. As the convergent tendency of modern inquiry became widely appreciated, immediate suspicions of scientific theft or fraud gave way to the resolution that independent multiple discovery is the default state of the scientific process—just a fact of scientific life.

A decline in priority disputes was not the only effect wrought by practitioners' knowledge of the convergent nature of scientific inquiry. Far more important than the incidence rate of priority disputes is the change it affected in scientists' approach to choosing a research problem (Merton 1961, 480). Because they know that multiple practitioners can and very frequently do make the same discovery independently, knowledge that someone is working on a given problem can dissuade a scientist from tackling that problem herself. Indeed, D. Keith Simonton, who conducted the first quantitative studies of multiples (Simonton 1979), remarked that "scientists will seldom work on an idea if they know that other scientists have already gotten a head start" (personal communication, April 29, 2014). It is worth reflecting for a moment on how much insight this fact gives us into the psychology of inquiry. Scientists are discouraged from working on a problem that others have already taken up because they take there to be something about the general nature of scientific investigation that makes discovery more or less inevitable; if a researcher begins a line of inquiry into a particular problem, eventually, she will solve that problem. If she started before you, she's probably going to solve the problem before you; something

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

exogenous to science would have to occur for her progress toward the inevitable discovery to be impeded. Closely related to this phenomenon is the not uncommon circumstance wherein a researcher will derive a result and decline to publish it because she assumes that it has already been discovered elsewhere—that her discovery is in fact a *re*discovery.

Granted, not all domains of inquiry are equally susceptible to independent multiples; it is questionable whether Simonton's remark captures contemporary pure mathematics and theoretical physics with as much accuracy as it does, say, microbiology. There are and will probably always be certain domains of inquiry that can benefit disproportionately from the contribution of unique degrees of human intelligence (more on this in section 4.6). In these domains, a late start can be compensated for by sufficiently distinctive ingenuity, and so practitioners working in these areas may be less prone to dissuasion than in domains that are more rapidly and reliably convergent. But a great deal of scientific inquiry simply involves what Kuhn called "mop-up" work—filling in well-defined gaps in a reasonably comprehensive image of nature. The results required to fill in these well-defined gaps are, in Bacon's memorable phrase, frequently "a birth of time, rather than wit" (quoted in Merton 1961, 473).

The last point about multiples worth mentioning here concerns their distribution across time as a function of the rate of scientific communication. Simonton found that as the rate of scientific communication increases, multiples tend to get closer together in time; given the nearly instantaneous rate at which contemporary scientific results are communicated, multiples are virtually simultaneous. This suggests two items of critical importance for appreciating the evolutionary nature of scientific inquiry. First, it suggests that, for a given line of inquiry, there are points at which a particular next step in the research process appears independently to multiple practitioners to be the most, or perhaps even the only, profitable direction in which to proceed. This is similar to the situation in which experienced chess players routinely find themselves, where a particular move stands out to competitors and observers as uniquely salient because of its clear superiority, even though multiple moves are technically available. The common link between these two phenomena and the adaptive process is that options for improvement tend to narrow as one approaches any optimum.

Second, it suggests that, in many cases, inquiry is naturally pushed to these points by previous achievements within the relevant disciplinary

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

community, and that the diffusion of scientific knowledge strongly determines how quickly these points are reached. Take a practitioner's acquisition of a given piece of scientific knowledge to represent his occupation of a particular point *A* in practice space; take the entire community's acquisition of a given piece of scientific knowledge to represent the convergence of each practitioner at point *A*; and take multiple discovery to be the movement of two or more community members to some point *A'* in practice space. As the rate of scientific communication across a discipline increases, the time at which community members acquire a given piece of scientific knowledge moves toward simultaneity. And as the rate of scientific communication increases, instances of multiple discovery likewise approach simultaneity. Expressed in terms of movement through practice space, this means that as the arrival of community members at a point *A* in practice space approaches simultaneity, the movement of two or more community members to some point *A'* also approaches simultaneity. As members of the same community, we can assume that those who arrived at *A'* moved there from *A*.

Now, if occupying point *A* in practice space did not increase the probability of occupying point *A'*, then there should be no reason that movement to *A'* should become simultaneous with the simultaneous convergence of practitioners at *A*; in other words, under those conditions, the size of the time interval between instances of a given discovery should be unrelated to whether the entire community has acquired a given piece of scientific knowledge. But it very clearly is related: The size of the time interval between instances of a given discovery decreases with the increasing rate of scientific communication precisely because the more practitioners possess a given piece of scientific knowledge, the more likely multiple practitioners are to pursue the same question and, eventually, draw the same inference. In the context of a given line of inquiry, the acquisition of more and more information about the relevant natural systems places more and more constraints on which future investigations might be profitable. At some point, the options for further inquiry become so limited that practitioners are, for epistemic and practical purposes, forced to pursue the same particular investigation. When every member of the community is simultaneously constrained with respect to future profitable investigations, it is inevitable that multiple practitioners will simultaneously conclude a particular investigation, drawing similar inferences.

I interpret Merton's idea that each scientific discovery is an "in principle multiple" to mean that inherent within the modern scientific approach to

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

inquiry is a strong directional tendency. This tendency puts serious constraints on how one can profitably move through practice space. Eventually, the constraints become so severe that only one movement proves to be profitable. Under these conditions, we naturally find multiple practitioners moving to the same point.

The directionality of modern scientific inquiry explains why practice space is so densely clustered: movement in systems that exhibit strong directionality is restricted and becomes increasingly restricted with time. The previous chapters have given us a couple of tools for assessing the degree to which scientific practice is susceptible to directional refinement through natural selection. We would expect highly susceptible systems to exhibit dense clustering in state space, and we would expect that clustering to be due, ultimately, to the convergence forced on them by the tightening constraints on profitable movement. And so it often is with scientific practice.

4.3 Adaptive Walks on Darwinian Landscapes

Even if convergence is symptomatic of a system under the influence of directional pressure, however, that does not thereby mean that the source of the directional pressure is natural selection. Many kinds of systems exhibit convergence due to the directional pressures that define those systems' dynamics; in only some of these systems is the convergence due to evolutionary adaptation. Rampant convergence within scientific practice was sufficient to bolster the claim that, in some cases, scientific practice is under directional pressure. But we're going to need something rather more specific in order to determine whether and to what degree the directionality exhibited by scientific practice is due to the fact that scientific practice is under the influence of the same forces that optimize phenotypes. We now want to know what it might look like for directionality of the scientific process to exhibit the signature properties of evolutionary adaptation—that is, what it would look like for the march toward points of convergence (or "in principle" convergence) in practice space to unfold in a characteristically Darwinian fashion.

To meet this demand, we turn to the specific sequence of changes that characterize an *adaptive walk*, a population's movement through morphospace toward a local optimum.[2] We can imagine, as an idealization, a population in a selective environment that never changes (indeed, laboratories routinely design experiments to approximate such a situation). As it

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

copies itself, the population's phenotypic composition shifts gradually as it increases its mean adaptedness. The limit of this process is one in which (1) the population's phenotypic mean is at a local optimum, (2) there is no accessible phenotypic state that improves adaptedness, and (3) even slight deviations from this optimal phenotype are penalized by selection. Now, of course, actual adaptive scenarios will not behave strictly in this way—for one thing, natural environments are not perfectly stable. But this is the canonical Modern Synthetic picture; adaptive scenarios in nature will be variations on this core theme. At its core, natural selection is the refinement of a population's morphology through the "preservation of favorable variations and the rejection of injurious variations" relative to some set of environmental pressures (Darwin 1859, 81). The variations in question can be thought of as deviations of arbitrary magnitude and direction relative to some population mean. For a given set of selective challenges, some variations are more favorable than others, and the population mean will continue refining in the direction of increasingly favorable variations until either (1) there is no variation that is more favorable, or (2) there is no more variation. Favorable variations are anything that improves an organism's ability to survive and reproduce. The classical image—Darwin's image—of these variations is that of "slight modifications" in morphology in a given direction—whatever direction causes an increase in adaptedness.

The magnitude of the modifications that compose an adaptive walk is of particular significance when considering what distinguishes Darwinian directionality from other forms of directionality. Faithful to Darwin's original impulse, the Modern Synthetic picture of adaptive modifications had been one in which a sequence of infinitesimally small changes led eventually to the ascent of a population to a local optimum. In Fisher's (1930) geometric model of an adaptive walk, the smallest favorable variation always possesses the highest probability of being favorable; as a variation's size increases, the chance that it might be favorable declines. For about 50 years, this result seemed sufficiently clear to support the conclusion that Darwin's adherence to the dictum *natura non facit saltum* had been well advised. Nature does not make a leap. Indeed, it moves as little as possible.

But as Kimura (1983) pointed out, Fisher's approach does not consider how the selective advantage of a given variation affects its probability of fixation. True, small variations are more likely to be favorable, but small favorable variations are also less likely than large favorable variations to have large effects on fitness. When the effect of selective advantage is

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

factored in, the magnitude of phenotypic changes composing an adaptive walk should have an intermediate value—not so small as to be selectively insignificant, but not so big as to be almost certainly lethal.

In a series of studies, H. Allen Orr (1998, 2000, 2002, 2003, 2006) introduced further refinements, which together give us the most specific picture to date of the sequence of changes that characterize adaptive walks up to a local optimum. Orr's fundamental insight (inspired by Fisher 1930, Kimura 1983, and Gillespie 1984) was that the size of a single adaptive modification tends to shrink over the course of an adaptive walk. This is due to the fact that the sphere of favorable variation in any direction shrinks as a population approaches an optimum. As such, the closer the population is to an optimum, the less likely a large variation is to be favorable. This trend is clearly reflected in the properties of Fisher's original geometrical model, although he seems to have been unaware of it (Orr 2005b, 12).

In addition, we are also able to say something about the magnitude of particular steps' contribution to the overall fitness gain made during an entire adaptive walk. Gillespie (1984) explicitly treated the adaptive walk as a particular type of sampling process, one that draws variants from the extreme right tail of a distribution. The logic underlying this approach is reflected in Fisher's model, principally in the fact that opportunities to increase fitness decline both in frequency and magnitude as the population approaches a phenotypic optimum. Using this sampling approach, Orr (2002) made three points of particular importance to our understanding of the process of adaptation, as well as to our understanding of how that process relates to the developmental sequence that characterizes changes in scientific practice. The first point is that the inaugural step in an adaptive walk accounts for, on average, a large portion of the total increase in fitness achieved during the ascent to a local optimum. The second point is that the largest single-step increase in fitness during an adaptive walk is, on average, at least 50% of the overall gain made during ascent. And third, Orr showed that the size of the fitness contribution of a particular adaptive modification declines as an approximate geometric sequence (i.e., any term $n+1$ in the sequence is a constant proportion of previous term n; see figure 4.1).

4.4 Epistemic Walks

I now want to make a number of points concerning how our understanding of the adaptive walk sheds light on what we know about the historical

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

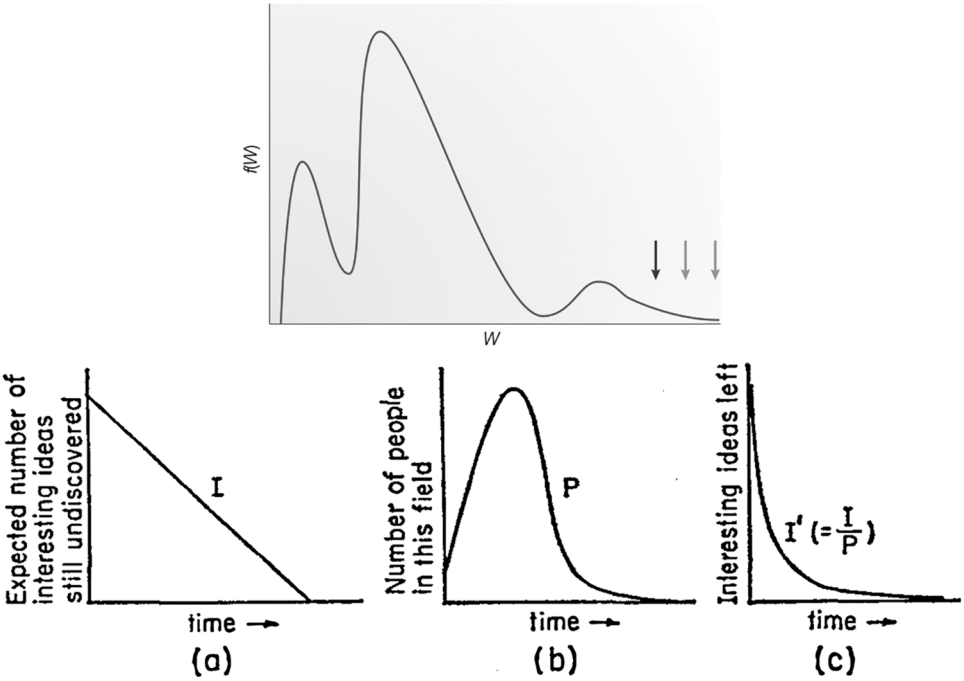

Figure 4.1
Orr's (2005b) distribution of fitnesses (*top*). Compare with Holton's (1962) "zeroeth-order approximation for a model of research in a specified area" (*bottom*).

development of scientific practice in certain cases. I first draw parallels between the basic pattern of change described by an adaptive walk on the one hand and a certain characteristic pattern of scientific change, moving from a research field's early days to its dotage, on the other. I then describe how we might understand the correspondence between that pattern of scientific change and the general pattern of adaptive evolutionary change in terms of similar forces are acting on both systems.

Let us begin with the point that a large percentage of the total increase in fitness achieved during an adaptive walk is due to the first step. What is the reason for this signature feature of adaptive walks, and in what way does it resemble the nature of changes in scientific practice in some cases? In biological contexts, we imagine that a phenotype begins its adaptive walk at some distance r from the phenotypic optimum; r is as far from the phenotypic optimum as one can be and still be viable in this population.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

At distances as great as *r*, evolution will favor relatively large single-step changes in morphology. This is because selection favors big fitness effects; in populations with high "Continuity," big fitness effects tend to only be possible through big changes in morphology. Thus, we should expect the magnitude of early changes to be very large in comparison with later changes. Orr's results prove this intuitive point (figure 4.2).

Given that the largest step in the adaptive walk averages around 50% of the total fitness gain, we can also infer that it too must occur in the relatively early stages of ascent to the optimum. For, if 50% of the total fitness gain is represented by *r*/2, a step of that magnitude becomes less and less

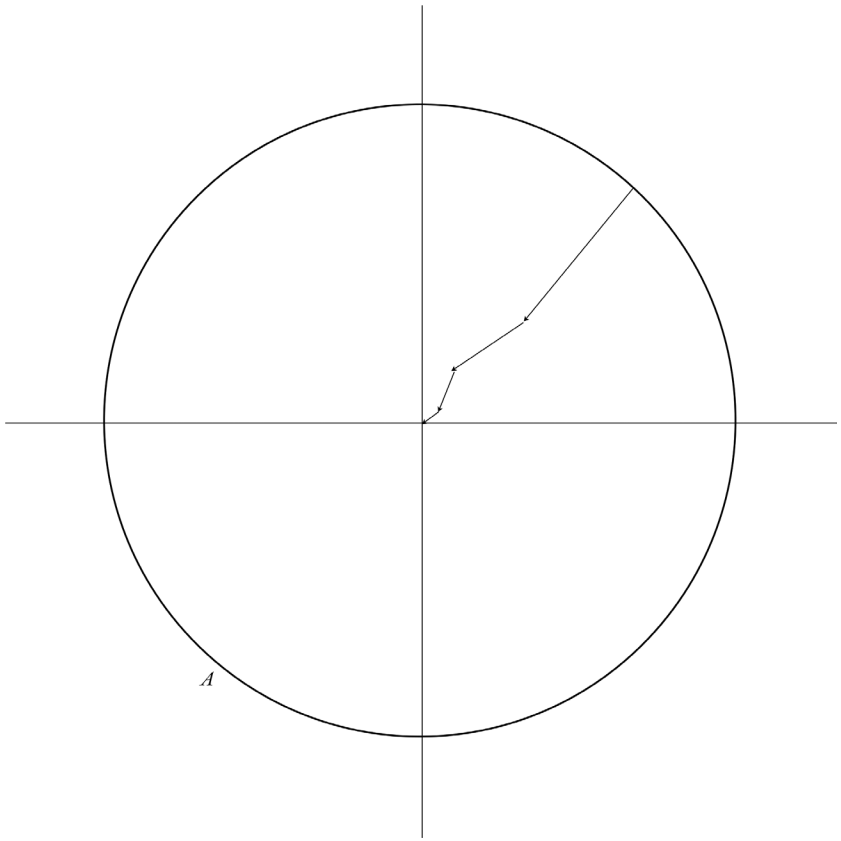

Figure 4.2
A sample walk: successively smaller changes reflect the shrinking supply of beneficial variations as the population approaches the optimum (from Orr 1998).

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

probable the closer we get to the optimum; smaller steps become equally profitable, and evolution will favor them (figures 4.3 and 4.4). This last point brings us to Orr's finding that the fitness contribution of a particular adaptive modification declines as an approximately geometric sequence.

Because each step is some fraction of the previous steps, the magnitude of steps late in the sequence relative to early ones is very small. Near the optimum, only small steps are profitable because anything too large will

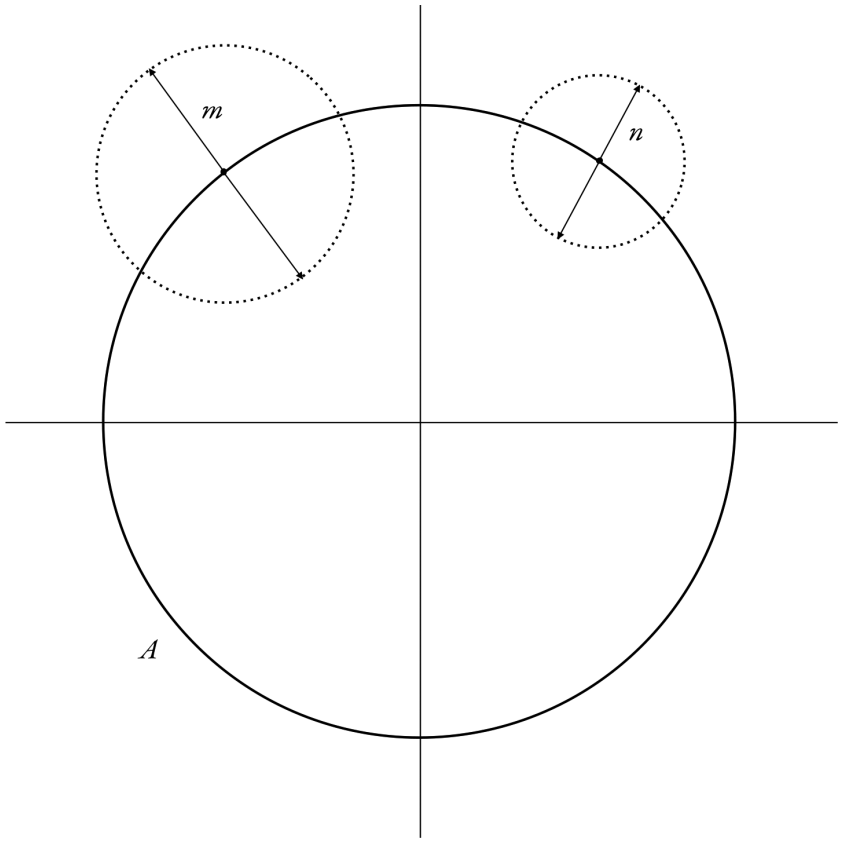

Figure 4.3
Upper limit: the range of evolvability (solid circle *A*) can accommodate 50% of variations of magnitude *m*, or much lower magnitude *n* (dotted circles); magnitudes denote the circle's radius. Beneficial variation (inside *A*) is as probable as deleterious (outside *A*). The phenotypic optimum is at the origin.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

overshoot the optimum and result in lower fitness, and the size of profitable steps gets smaller and smaller as the optimum is approached.

If we are inclined to view the adaptive walk in paradigmatically Darwinian populations as a model of (some aspects of) the scientific process in idealized scientific communities, the significance of these results is threefold. First, they suggest that the overwhelming majority of progress during a particular episode of scientific advance can be credited to one or two major contributions, while the remaining progress differential is made up through a relatively large number of small contributions. Second, they suggest that these major contributions occur early in the course of a particular

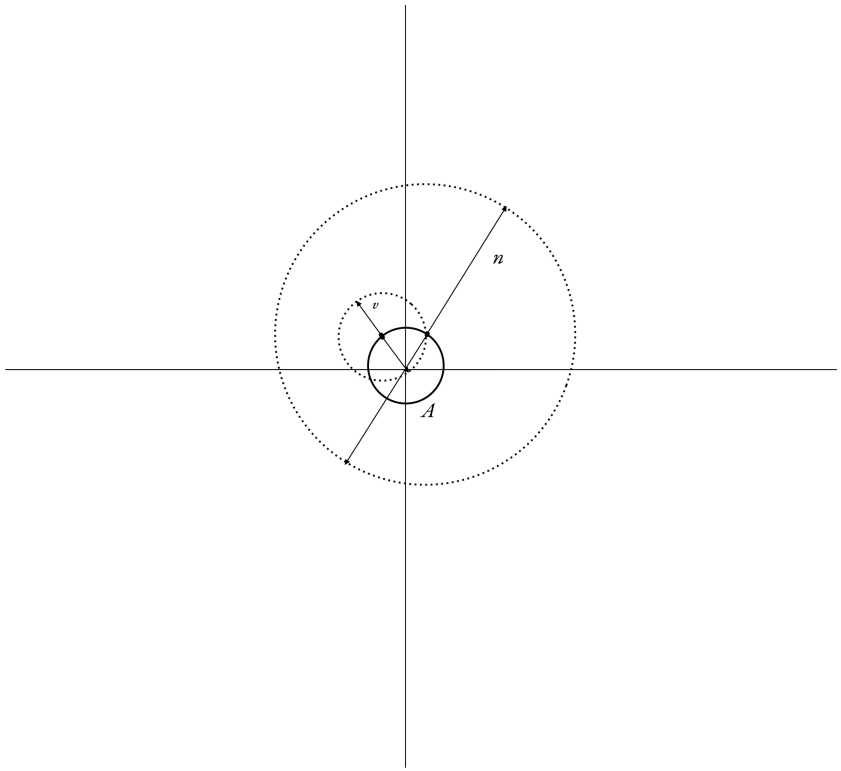

Figure 4.4
Lower limit: the range of evolvability (*A*) near the optimum. Unlike what is shown in figure 4.3, *n*-sized variations have only a small chance of evolving, but lower-magnitude *v*-sized variations have a higher probability.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

episode. Together, they point to a familiar feature of the historical pattern that characterizes scientific advance—namely, that of one or two early, monumental achievements, followed by a relatively protracted period of smaller, related contributions. As Kuhn famously described the historical role of the paradigm-inducing classic works of science,

> these and many other works served for a time implicitly to define the legitimate problems and methods of a research field for succeeding generations of practitioners. They were able to do so because they shared two essential characteristics. Their achievement was sufficiently unprecedented to attract an enduring group of adherents away from competing modes of scientific activity. Simultaneously, it was sufficiently open-ended to leave all sorts of problems for the redefined group of practitioners to resolve. (Kuhn 1962b, 10)

The large early steps that characterize adaptive walks can be viewed as achievements "sufficiently unprecedented to attract an enduring group of adherents away from competing modes of scientific activity." The reason for their being "sufficiently unprecedented" is obvious enough: Orr's model shows that any large steps must come on the heels of a protracted series of small modifications. The works cited by Kuhn—Franklin's *Electricity*, Newton's *Principia* and *Optics*, and Lavoisier's *Chemistry*—were unprecedented in precisely the sense that no single achievement had contributed a remotely comparable degree of scientific advance in a very long time, if ever. In each case, for example, the preceding century had hosted a number of important results, but in no instance do we find a contribution whose magnitude approaches that of these age-defining classics of science.

The characteristics of the adaptive walk, then, give a clear meaning to the idea that major contributions to science tend to be "sufficiently unprecedented." But the concept of an adaptive walk is at its most luminous when aimed at understanding how the effect of these major contributions is to "attract an enduring group of adherents away from competing modes of scientific activity." This brings us to the third point. Let us conceive of a range of evolvability as the circle with phenotypic optimum at center *O*, where the population currently resides on the edge and where each point on or within the circle represents an evolutionarily possible phenotypic state. The states are "evolutionarily possible" in the sense that they are no farther from the optimum than the population's current state, and thus allow a population to travel to any one of them through well-understood evolutionary mechanisms.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Imagine this population sitting at a point in morphospace that is equidistant from several optima (figure 4.5). Suppose that after phenotypic change *p1*, the population has traveled some large percentage of the distance to optimum O_p.[3] In the ideal case of nonoverlapping evolvability ranges associated with these optima (figure 4.6), at *p1* the population has become closer to O_p than to any other optimum (where the probability of remaining equidistant after a change varies inversely with the size of the change, and directly with the degree of overlap between adaptive spaces).[4] At that point, travel to optima other than O_p will be precluded, because any such travel would reduce the population's mean fitness. Members of the population varying in the direction of O_p will be rewarded with higher fitness, whereas those varying toward optima O_q, O_r, and O_s, will be penalized.

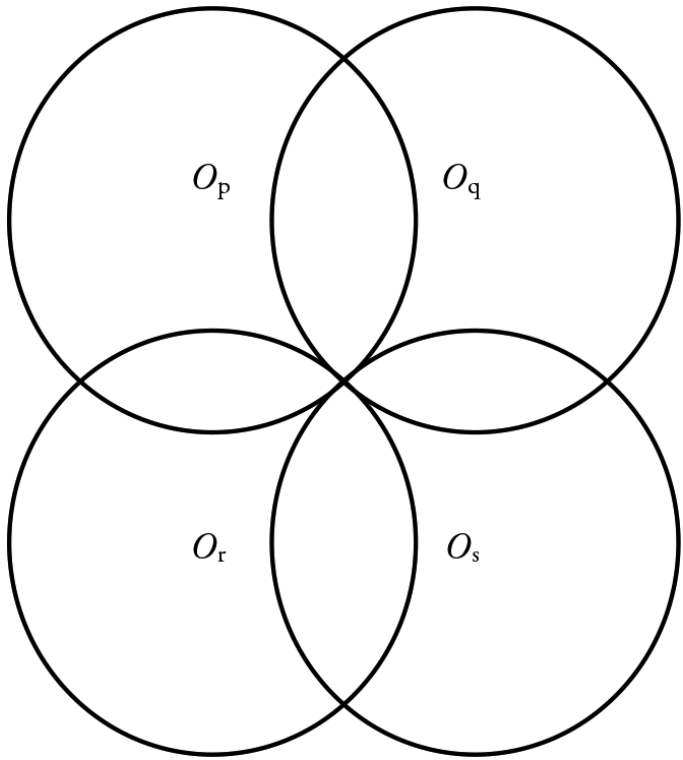

Figure 4.5
The population at the center, equidistant from four distinct morphological optima.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

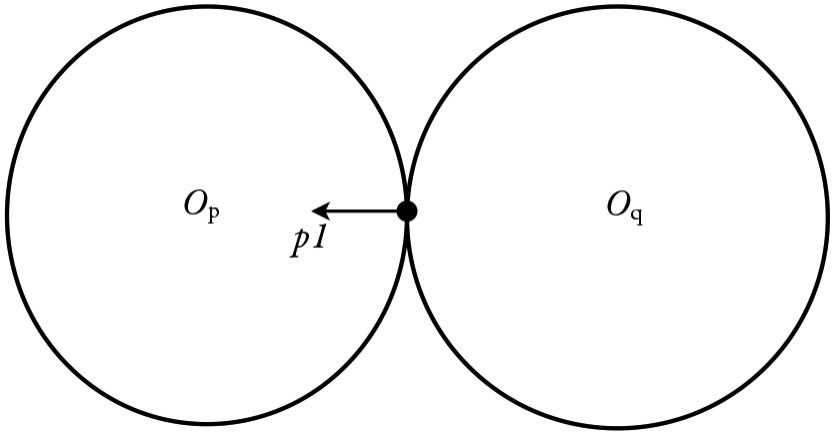

Figure 4.6
The population has moved toward a particular optimum.

It is in this sense that a large initial step "attracts an enduring group of adherents away from" travel toward alternative optima: once the population has been carried much of the way to a fitness optimum, the fitness costs of pursuing alternative phenotypic strategies are unsustainably high. Smaller initial steps will be less successful at precluding subsequent travel toward other optima, for two reasons. First, the small selective advantage associated with small steps makes them susceptible to loss through random drift. Second, it is much easier to change course after a small step than after a big one; the smaller the step, the less one gives up in abandoning it. The less one gives up, the less one needs to compensate, and small compensatory gains are much more plentiful than large ones.

How does this picture relate to the effects of monumental contributions to scientific knowledge? Let us begin with Kuhn's image of "pre-paradigm," which states:

> In the absence of a paradigm or some candidate for a paradigm, all of the facts that could possibly pertain to the development of a given science are likely to seem equally relevant. As a result, fact-gathering is a far more nearly random activity than the one that subsequent scientific development makes familiar. (Kuhn 1962b, 15)

Kuhn's favorite example of a pre-paradigm state was optics before Newton, which he evocatively described as "something less than science":

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> Being able to take no common body of belief for granted, each writer on physical optics felt forced to build his field anew from its foundations. In doing so, his choice of supporting observation and experiment was relatively free, for there was no standard set of methods or of phenomena that every optical writer felt forced to employ and explain. (13)

Now, Kuhn was careful to emphasize that inquiry is by no means aimless under these conditions; there are scores of focused, well-executed empirical and theoretical studies to which we could point, from Al-Hazen to Hooke. But, as he says, "though the field's practitioners were scientists, the net result of their activity was something less than science" (1962b, 13). The period prior to Newton was not wanting for important results relevant to our understand of light's behavior. But, taken together, those results did not amount to scientific progress in any meaningful sense, because there was no generalized movement of the group of practitioners through practice space.

We can relate this to the Fisher-Orr geometrical model in the following way. Kuhn's pre-paradigm state can be represented as a one-dimensional practice space with no optimum—say, a straight line on which all of the field's practitioners are located (figure 4.7). The one-dimensional nature of practice space under these conditions reflects the fact that, although practitioners can be more or less similar to one another (represented by the intervals between them), there is no set of fixed points in practice space relative to which we could rank each practitioner's distance—no "standard set of methods or of phenomena that every optical writer felt forced to employ and explain" (Kuhn 1962b, 13). No wonder that in these states, inquiry often consisted largely of documenting the differences between the fundamental commitments of one's school of thought (one's "particular metaphysic") and those of his opponents.

For Kuhn, the publication of Newton's *Opticks* (or his *Principia*, or Darwin's *Origin*, or whatever) imposed a kind of structure on practice space, such that all practitioners could locate their position in it relative to both (1) where they had previously been and (2) where they needed to go next. In each case, the imposition of this structure was made possible by the author's development of a comprehensive approach to different kinds of

Figure 4.7
A pre-paradigm practice space.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

problems in the field, combined with illustrations of the new approach's power in the form of reasonably successful applications. The success of these applications allowed practitioners to see, retrospectively, how far this giant leap in scientific understanding had taken them. The new perspective on inquiry afforded by something like Darwin's *Origin*, for example, gave even leading naturalists like Joseph Hooker a vivid glimpse of how unenlightened their previous studies had been:

> What Darwin constantly emphasized, and what his contemporaries recognized, was that the *Origin* was not only a confession of ignorance but also a structuring of our ignorance. . . . [I]ts primary accomplishment lay in identifying the questions that biologists ought to ask. It is because of this primary accomplishment that Darwin may truly be said to have revolutionized the field. The nature of that revolution is captured in one of Hooker's letters to Darwin:
>
> > But, oh Lord!, how little do we know and have known to be so advanced in knowledge by one theory. If we thought ourselves knowing dogs before you revealed Natural Selection, what d—d ignorant ones we must surely be now we do know that law. (Quoted in Kitcher 1985, 60)

At the same time, the generality of the approach, combined with the form and substance of the problems to which it was applied, pointed to what sort of future research would now be possible.[5]

In terms of the Fisher-Orr model, the publication of these paradigm-inducing contributions allowed for the emergence of an optimum in practice space, however vaguely defined at first. Practitioners were thereafter able to gain a rough sense of where their newly developed level of understanding stood in relation to what their previous state had been, as well as how relatively unenlightened were those individuals who failed to take advantage of the new developments and thus remained in what was effectively a pre-paradigm state. More importantly, practitioners were able to gain a rough sense of where their newly developed level of understanding stood in relation to what it could eventually become. Initially, and for some time afterward, the practice optimum would persist as an amorphous cloud of points, which becomes increasingly well defined as the actual historical course of inquiry unfolds (figures 4.8 and 4.9).

As in the standard adaptive walk, the emergence of an optimum (or an optimal neighborhood) in practice space results in the immediate exclusion of many forms of inquiry as insufficiently fruitful to warrant further pursuit. Subsequent to these early foundational events, research proceeds

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

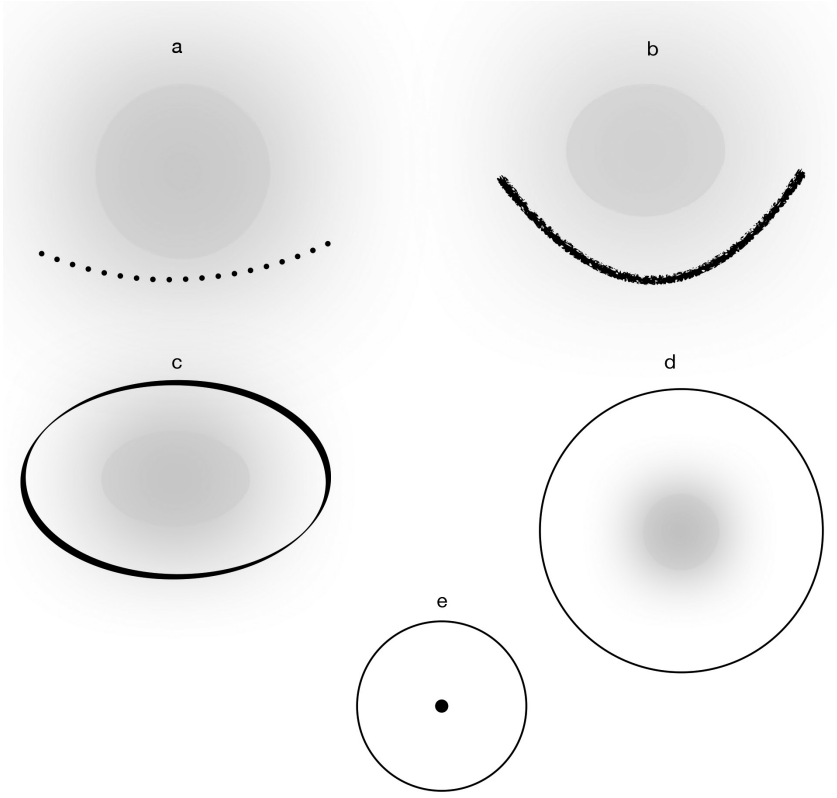

Figure 4.8
The practice optimum persists as an amorphous cloud of points, gradually gaining definition as the actual historical course of inquiry develops.

largely as attempts to follow up on the largess of problems bequeathed to the newly defined community in light of their adoption of the new framework. Kuhn dubbed this period one of "normal science," which he vividly taxonomized as an array of different kinds of research aimed at the further investigation and articulation of quantities and concepts that, subsequent to paradigm adoption, have been shown to be of particular significance for the understanding of nature. According to Kuhn's picture, the high magnitudes of the early, foundational contributions stand in contrast to the size of the contribution made by these various bits of "mop-up work," which tend to be relatively small and tend to get smaller and smaller as research develops within a particular framework. A related point was later made by

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

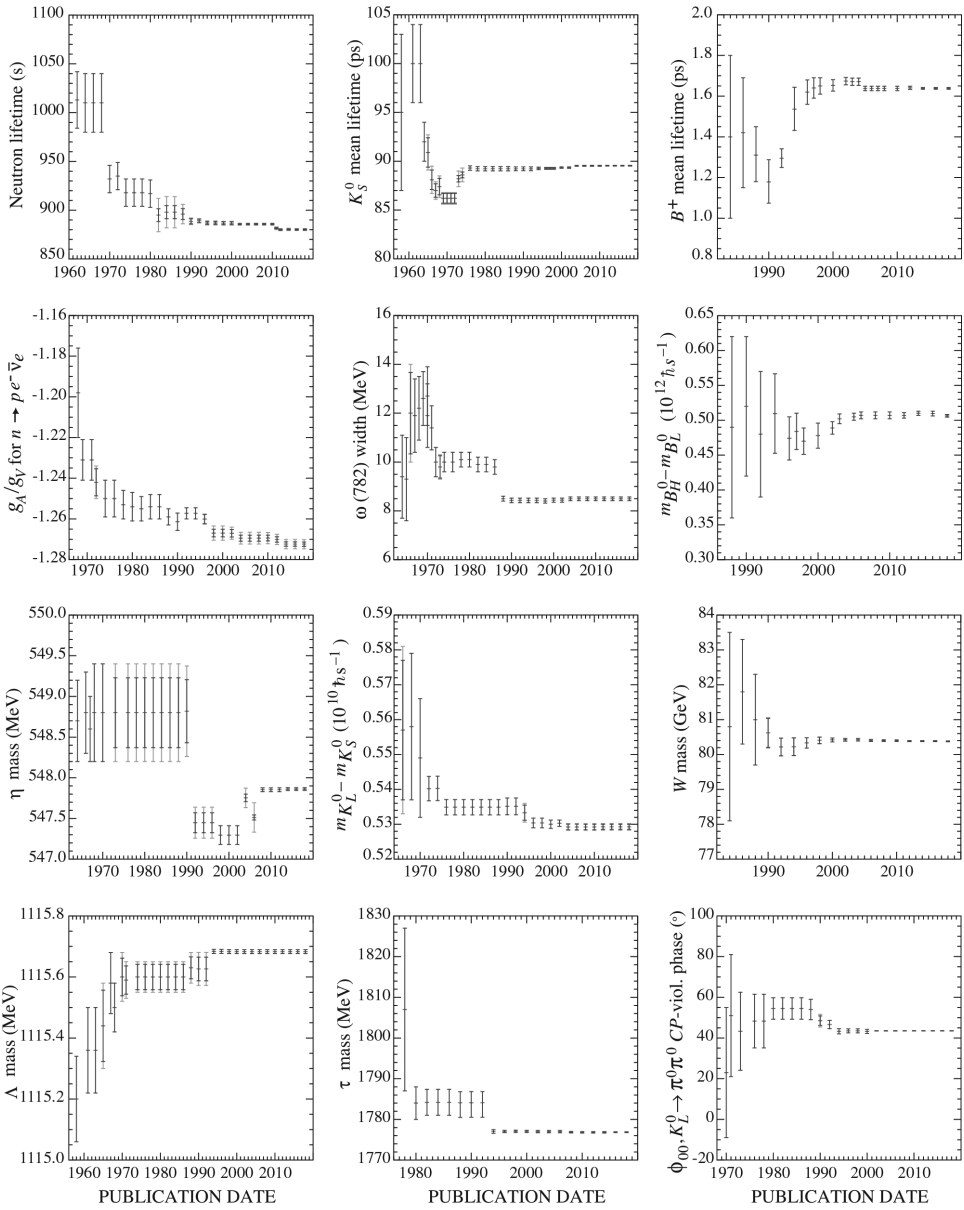

Figure 4.9

From the *Review of Particle Physics*: "A historical perspective of values of a few particle properties tabulated in this Review as a function of date of publication of the Review. A full error bar indicates the quoted error; a thick-lined portion indicates the same but without the 'scale factor'" (Zyla et al. 2020, 19). The asymptotic progression of the ranges of values mirrors the progression described by Orr and hypothesized by Holton in his "zeroeth-order approximation." See figure 4.1.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Mulkay (1975, 6–7), who observed that "research areas tend to develop in response to major innovations which appear early in the growth sequence, such as that of Mendel in the case of genetics. Subsequent work tends to consist primarily of elaborations upon these central contributions." This geometric decline in magnitude is precisely the one Orr found to be a stable feature of adaptive walks.

In plainest terms, what accounts for the tendency of research contributions in these episodes to decline in magnitude is the basic fact that researchers working within a scientific community are preoccupied with building on the community's previously established modifications to scientific practice. Whether the modifications be with respect to theory, methodology, or what have you, all new results must be reconciled with the community's existing range of consensus. As that range of consensus expands, it is increasingly likely that a major change to scientific practice will conflict with an element of practice on which the community has already achieved consensus. Accepting a major change to scientific practice during later stages of paradigm-governed research will require the community to divest itself of an item of existing consensus. And given that research builds on previous research, a sufficiently large change could ultimately require major revisions to a large body of existing scientific knowledge. This option is typically a last resort. Well-functioning scientific communities seem to prefer smaller modifications that are consistent with existing knowledge, where they are available. We should expect major changes to scientific practice only when the high costs of doing so are outweighed by the benefits. We have a clear and direct explanation for this phenomenon from the perspective of the Fisher-Orr model of an adaptive walk. We will examine instructive counterexamples to this tendency toward smaller and smaller modifications in chapter 6, when we discuss the kinds of major changes to scientific practice that constitute major scientific revolutions.

In idealized scientific communities, the space of potentially fruitful changes to practice shrinks to the point where only one move can generate further profit—namely, the move to the optimum state. These are the stages in the development of scientific practice where we expect to see the kinds of multiple discoveries that Merton and others have described. We can thus see how we might conceive of frequent instances of convergence in practice space as unfolding in a manner that corresponds to some version of a paradigmatically Darwinian population's adaptive walk. The significant point

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

is not merely that different scientists working independently often hit on the same idea. Nor is it that the nature of scientific inquiry itself is such that these multiple discoveries are inevitable. Rather, the significant point is that, in these cases, multiple discoveries are inevitable when scientific inquiry is governed by the same forces that produce highly refined adaptations in biological organisms.

The pattern described above, in which a large single modification initiates a directional trend characterized by a subsequent series of increasingly smaller modifications, has been shown by Orr to be a stable feature of adaptive walks in general. I believe it is no accident that the same pattern was identified by Kuhn as the characteristic arc of scientific development. Both patterns are generated by the adaptive process. Even more interesting, I think, is the way Kuhn tends to fixate on precisely those phenomena that are most significant from an evolutionary perspective. One critical question that our preceding discussion leaves open is how exactly the adoption of a paradigm facilitates the emergence of a practice optimum and provides a structure to practice space in the manner described above. In Kuhn's view, this structure comes not so much from any particularly brilliant insight into the secrets of nature, but *merely from the agreement of a sufficiently large number of individuals on a sufficiently wide range of issues*. Nowhere is he more clear on this point than when he observes that "it is sometimes *just its reception of a paradigm* that transforms a group previously interested merely in the study of nature into a profession or, at least, a discipline" (Kuhn 1962, 19; emphasis added): "Men whose research is based on shared paradigms are committed to the same rules and standards for scientific practice. *That commitment and the apparent consensus it produces are prerequisites for normal science, i.e., for the genesis and continuation of a particular research tradition"* (11; emphasis added).

In other words, the actual substance of the group's agreement—their specific shared commitments—was, for Kuhn, not particularly important for explaining why science seems to be inherently progressive. Rather, the causally significant aspect of paradigm adoption is agreement per se. "Commitment and the apparent consensus it produces are prerequisites for normal science" in the sense that consensus shuts down the interminable foundational disputes that characterized pre-paradigm inquiry, allowing research to build on previously certified results. The fact that there is now widespread agreement on *something* means that inquiry can proceed

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

to more refined states of investigation. In evolutionary terms, consensus constitutes the fixation of certain points in practice space. These fixed features of practice are analogous to the stabilized environments described in chapters 2 and 3. We noted there that stable environments represent the constraints that define optima in a state space. In Kuhn's picture (suitably attenuated), we do not quite get something like *science* in pre-paradigm states, because there is no stable environment that would allow inquiry to embark on the adaptive walk that characterizes the gross pattern of scientific development.

4.5 Directionality in the Small

Orr's studies of the adaptive walk revealed a fourth general aspect of adaptive walks that is relevant to our understanding of the directionality of science, one that reflects the fact that, like many historical phenomena, the development of scientific knowledge exhibits important regularities at multiple scales. In particular, Orr argued that the pattern described above—a big change followed by a geometrically decreasing series of smaller and smaller changes—is something that characterizes the adaptive walk at each point on its trajectory. Whether we are looking at the refinement of a trait from, say, (1) the period spanning its first appearance in the population to its eventual fixation and optimization or (2) the period between near-optimality and optimization, the same pattern will arise. This robustness across scales is due to the fact that, as Orr (2005b, 12) pointed out, "adaptation is characterized by a kind of repeated re-scaling. . . . [T]he population confronts essentially the same problem at each substitution, but on a smaller scale." Once again, the geometrical properties of Fisher's model provide an intuitive representation of this pattern and its robustness across scales. After each step (in which the adaptive walk is redefined as a walk across some distance r' [$< r$]), the points within the new circle retain the same properties relative to one another and to the optimum that they had in the previous step; changes that are large relative to r' (rather than r) will still be favored early on, for the same reasons. Thus, the size of early changes relative to the overall distance to the optimum remains constant throughout the adaptive walk.

If this is correct, we should expect small-scale phenomena like the pursuit of a single line of inquiry during a portion of one scientist's career

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

to exhibit the same basic pattern of adaptive change characterized by very large-scale changes, such as the transition from Galileo's pre-classical mechanics to Lagrange's perfection of classical mechanics roughly 200 years later. I now aim to show that the refinement of scientific practice on scales much smaller than these major transitions often proceeds in this manner. In the remainder of this section I describe several different kinds of episodes in a scientist's career that begin with a relatively high-magnitude change, followed by a series of variations and elaborations that grow smaller as they near some well-defined point in practice space. I will call any such relatively high-magnitude change a *foothold*.

A scientist's career, or at least a major chunk of it, is occasionally defined by a single event that occurs at an early stage in her professional trajectory. Scientists frequently describe the period leading up to such an event as consisting of much searching, false starts, blind alleys, and dead ends, making essentially no progress despite an enormous amount of effort. At some point, one is given a brief but highly resonant glimpse of a way forward, and thereby gains a foothold on a path that will eventuate in scientific achievement. The years that follow consist primarily of a campaign to bridge the now-perceptible gap between one's foothold position and the achievement to which it will eventually lead.

Certainly the concept of a foothold is exemplified by the well-known "flash of brilliance." But it's important not to invest brilliance flashes with too much influence over how knowledge grows. The overwhelming majority of individual contributions to scientific knowledge are not connected with brilliance in any way. And, at any rate, the sort of initial change I have in mind is generally much more modest in scale than whatever genuine brilliance seems capable of fostering (more on this in section 4.7). For example, one type of foothold consists merely in finding a really good research problem to work on. The search for a good problem can be an epoch-defining node in the career of a researcher, as whatever problem she settles on may have a strong influence over the next several years, if not decades. Indeed, having a nose for a good problem is so significant that it is one of the qualities that distinguishes the merely competent scientist from the truly great. As particle physicist Savas Dimopoulos observes,

> In particle physics, you have to have a threshold amount of intelligence, whatever that means. But the thing that differentiates scientists [from one another] is purely an artistic ability to discern what is a good idea, what is a beautiful idea,

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> what is worth spending time on, and, most importantly, what is a problem that is sufficiently interesting, yet sufficiently difficult that it hasn't yet been solved, but the time for solving it has come now.[6]

Once a research problem has been adopted, a new phase in the scientist's career has begun. This phase primarily involves breaking the research problem down into manageably sized subproblems and then going about the business of solving them. In order to make genuine progress in a sufficiently short period of time, she may be fairly restricted with respect to her ability to work on new problems or on problems that are not directly related to her previously chosen overarching research program. Here, again, we see a reflection of the general dynamics underlying the adaptive walk: a grand initial step, followed by relatively small movements whose reward is a function of (1) their size, and (2) how closely they remain on the straight line connecting the initial step to the optimum that emerges in response to it.

Another very common species of foothold is the sort whereby a practitioner is given a glimpse into a pathway to a solution for a theoretical problem. These episodes are often described as vague but convincing impressions of the final form that a solution will eventually take, but whose completion might be years away, if not longer. Darwin's reading of Malthus in 1838 is a classic example, this event having at last given Darwin a firm grip on the general means by which he might treat the major theoretical problems associated with species over which he had toiled since his voyage on the *HMS Beagle*. In Malthus's "principle of population," Darwin instantly recognized the outline of a mechanism that could explain how species become adapted to their environments. What followed over the next decade and a half was Darwin's attempts to fill in the empirical details that would allow his newly formulated mechanism to achieve the standard of a *vera causa* (Ospovat 1983). His effort in this direction is less apparent in the *Origin*, where the description of natural selection is highly schematic and aimed primarily at demonstrating the explanatory adequacy of the mechanism, as well as his role in its discovery. Rather, the gap-filling measures that follow an adaptive walk are more clearly reflected in Darwin's unfinished, detail-laden *Natural Selection*, whose composition was interrupted by A. R. Wallace's letter and the subsequent need to write up an abstract of his views (i.e., the *Origin*) (Stauffer 1975). Tappenden (2005) argues for the significance of this phenomenon in mathematics, describing a variety of contexts in which a fleeting, impressionistic glance of what form a proof will

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

eventually take generates a mental map for how to actually construct the details of the proof.

Further, perhaps more tangible manifestations of the "foothold" phenomenon come in the form of the different types of concrete conceptual and methodological innovations that we typically associate with scientific change. For example, practitioners often find themselves in the position of having a research problem of central importance to the field but with no apparent means of empirical resolution. Under these conditions, a practitioner's greatest potential contribution to the problem's solution will often come not in the form of empirical results that actually resolve the issue, but in the form of a methodological innovation that allows her (or whomever) to go about the quotidian business of collecting data that bear directly and decisively on the hitherto vexed problem. A particularly strong illustration of this variety of foothold is the invention of the telescope and Galileo's subsequent use of it to investigate (and confirm) many of the implications of the Copernican model (Swerdlow 1998). Many of the model's distinctive predictions simply could not be confirmed at distances that were visible to the naked eye. The major breakthrough was the invention of something that could, in principle, provide an empirical resolution to these problems. Here, the small subsequent steps characteristic of an adaptive walk came in the form of comparing, one by one, the model's predictions with what one was now capable of observing. This is not to downplay the difficulty and complexity of figuring out how to get the model to bear on nature, nor to denigrate the enormous ingenuity displayed by Galileo in his efforts to apply telescopic power to the problems of the planets. But the astronomical community already had some 60 years to work out the many distinguishing observational consequences of the Copernican model, and the phenomena to be expected were by 1609 very well understood. Similarly, actually checking the predictions, although a delicate and laborious business, can be accurately described as primarily "mop-up work" in Kuhn's sense. What moved the community from its stalled position after 1543 to a rapid resolution of many open questions was the single methodological step that dramatically improved their observational powers.

As a final species of foothold, consider the phenomenon whereby a conceptual innovation leads to the generation of a vast range of solvable problems. Naturally, in cases where the effect of the innovation is to generate new kinds of research problems, the consequences will tend not to be

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

confined to a single scientist's career. A classic illustration of the relevant sort comes from Dalton's conceptualizing elemental substances as integral multiples of individual atoms (and compounds as combinations of integral multiples of elements in regular proportions). Conceiving of substances in this way allowed for the relatively rapid development of an entire industry of empirical studies of the chemical composition of specific substances. Within less than half a century, the community of chemists had achieved consensus on chemical formulas for a wide variety of substances. Fundamentally, what made this rapid convergence possible was the assumption that an element, being as it was composed of distinct species of atom, can be distinguished by its weight relative to other elements. Chemical compositions could then be fairly easily and definitively investigated simply by measuring changes in weight (or in volume) that occurred through experimental separation of elements (e.g., through combustion). Interestingly, as Rocke (1984) shows, the consensus events themselves exhibit the signature pattern of the adaptive walk. Major, qualitative differences in chemists' views on various chemical formulas were overcome more or less instantaneously following Dalton's reconceptualization. What followed was the shoring up of different community members' estimates of atomic weights, which would proceed from relatively large disagreements (of 1/10 or more), to the vanishingly small differences made possible by the improvement and standardization of measurement techniques, such as those described in chapter 2.

The various types of footholds surveyed here all exhibit the big-to-small-to-smaller pattern of change now familiar from our discussion of the adaptive walk. True to Orr's insight, this pattern holds regardless of the timescale across which the focal episode of change is to have occurred. And his explanation for why the pattern is timescale invariant certainly makes sense in the context of scientific inquiry: "the population confronts the same problem . . . but on a smaller scale" (Orr 2005b, 12). The challenge that biological populations confront is to drive to fixation some variation that will maximize fitness in a given environment, faced with the unavoidable fact that most large variations will negatively impact fitness, and that smaller variations are more likely to be lost by chance unless they occur when the population is adjacent to the phenotypic optimum. Likewise, at each stage, whether of an individual scientist's career, or of the historical development of an entire research field, the progress of inquiry is confronted with the

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

same challenge. That challenge is to introduce some modification to scientific practice—be it different problems, different explanations, different methodologies, different technologies—that will generate new, yet closely related, lines of inquiry. Every iteration of inquiry confronts this challenge in the context of an analogous, timescale-invariant set of constraints—namely, (1) that almost all very big changes to practice will impede or reverse the progress of science by extinguishing existing lines of inquiry, and (2) that very small changes to practice are far less likely to have a lasting effect on inquiry unless they occur at a time in which a problem's resolution is nearly at hand. If approximation to a paradigmatically Darwinian population was not governing scientific change in these instances, there is no clear reason to expect this same pattern to crop up over and over again in the context of scientific inquiry. Why, for example, is scientific change in not frequently characterized by increasingly larger changes to practice? The fact that this bigger-to-smaller pattern dominates so many contexts is a puzzle—one that is solved by the notion that scientific practice is often significantly modified by natural selection on account of the fact that practice often significantly resembles a paradigmatically Darwinian population. Ceteris paribus, the degree to which the scientific community's practice proceeds in conformity with a canonical adaptive walk should vary directly with the degree to which that community's practice approximates a paradigmatically Darwinian population.

4.6 A Leap across the Landscape

In his effort to build the theory of natural selection on the most conservative foundation he could, Charles Darwin adopted a particularly austere view about the magnitude of natural variations:

> As natural selection acts solely by accumulating slight, successive, favourable variations, it can produce no great or sudden modification; it can act only by very short and slow steps. Hence the canon of "Natura non facit saltum," which every fresh addition to our knowledge tends to make more strictly correct, is on this theory simply intelligible. We can plainly see why nature is prodigal in variety, though niggard in innovation. (Darwin 1859, 471)

Unlike most of the empirical regularities on which Darwin relies in the *Origin* and in his later works, the notion that "nature does not make leaps" was not one that was widely embraced by his contemporaries. And despite

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

its enduring popularity among developmentally oriented practitioners throughout the nineteenth and twentieth centuries, it is only relatively recently (due largely to our ability to look closely at gene sequences) that the community of evolutionary biologists has come to accept the evolutionary significance of very large mutations. Contrary to what Darwin thought, nature does occasionally make leaps, and those leaps have a major effect on which point in morphospace emerges as the phenotypic optimum.

In a similar fashion, Merton (1961, 486) considered the proverbial "leap from antecedent knowledge" to be related to the directionality of science in a central, important, and fascinating way, one that connects it directly to our discussion of the characteristic pattern of adaptive change. He observed that

> the greatest men of science have been involved in a multiplicity of multiples. This is true for Galileo and Newton, for Faraday and Clerk Maxwell, for Hooke, Cavendish, and Stensen, for Gauss and Laplace, for Lavoisier, Priestley, and Scheele—in short, for all those whose place in the pantheon of science is beyond dispute. (484)

Merton argued that this fact suggests a conception of *scientific genius* according to which "the individual man of scientific genius is the functional equivalent of a considerable array of other scientists of varying degrees of talent" (484). For, what the individual genius is able to see entirely on his own in a "leap from antecedent knowledge" will frequently be rediscovered a generation or more down the road after an entire field (composed of nongeniuses) has reached the same destination through a protracted series of comparatively modest modifications to scientific practice. In Darwin's terms, the scientific genius is often able to achieve in a single stride "the great effect produced by the accumulation in one direction, during successive generations, of differences absolutely inappreciable by an uneducated eye" (Darwin 1859, 32).

At the beginning of this chapter I argued that the routine occurrence of multiple discoveries in the history of science can be understood as a reflection of the presence of a directional tendency inherent to scientific inquiry, for the same reason that morphological convergence is evidence for the directional influence of natural selection. What both phenomena indicate is that, in each of their respective domains, certain end-states are *historically robust*; had circumstances been different—had, for example, Isaac Newton never existed—the mathematical physics community would still have figured out that the elliptical orbit of the planets was caused by an attractive force emanating from the sun that diminished with the square of a planet's

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

distance from it, probably within a few short years of Newton's derivation (Westfall 1980; Cohen 1983). The reason these states are historically robust is because they are virtually inevitable, given (1) the range of possibilities and (2) the causal factors that significantly affect the historical development of scientific knowledge in these contexts.

Understood in Merton's terms, scientific genius offers a new kind of support for the idea that some points in practice space are historically robust, in that it seems to show that the tendency to occupy certain points is (to a degree) *path-independent*. That is, whether the community converges on certain points is not always dependent on whether its members have come to occupy adjacent points in practice space. Now certainly, Merton observes, "some [discoveries] flow directly from antecedent knowledge in the sense that they are widely visible implications of what has gone just before" (1961, 486). We often get multiple discoveries under such conditions precisely because a particular, relatively modest refinement to scientific practice stands out to a field's practitioners as the obvious next step. Very occasionally, though, that same point in practice space is salient only to one practitioner, and from a great distance—a distance so vast that it may not even appear coherent or relevant from the perspective of the rest of those working in the field. One cannot help but quote Schopenhauer here: "The man of talent is like the marksman who hits a mark the others cannot hit; the man of genius is like the marksman who hits a mark they cannot even see to" (Schopenhauer 1909, 158). Given enough time, the community of average practitioners will arrive at that point, but only because its contours and relevance have gradually become salient as a result of "numerous, successive, slight modifications" to scientific practice.

The great "leaps from antecedent knowledge" to ideas that are arrived at under other circumstances by small steps on the part of many individuals is a powerful illustration of how some ideas might function as attractors in practice space, just as phenotypic optima function as attractors in morphospace. Certain phenotypic states pull populations toward them because they maximize fitness under a relevant set of constraints. Analogously, I suggest that certain points in practice space are more or less inevitable because they maximize something akin to fitness in science: *fruitfulness*. Exactly what fruitfulness is and why its maximization is fundamental to the progress of science is not something I can delve into at this point. These questions are given careful consideration in other work (Haufe 2023).

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

4.7 Conclusion

Our discussion of the correspondence between the evolutionary process and the scientific process has now reached a critical juncture. The previous chapters endeavored to establish a number of highly suggestive analogies between the core factors involved in the process of evolutionary adaptation on the one hand and some of the social phenomena known to be causally influential in certain scientific communities on the other. If the analogy is more than suggestive—that is, if the resemblance of the two systems is explained by the fact that they are different instantiations of the same general causal process—the correspondence between these two systems should extend beyond a synchronic similarity between their causally relevant factors. We should also expect the historical patterns generated by the causal interaction of their respective factors to be essentially in agreement. Consequently, this chapter has argued that the general historical pattern of adaptive change frequently matches, in its most central respects, the historical development of scientific knowledge, in its most central respects.

As in many other places in this book, I have here derived my understanding of the basic pattern of historical development in science from certain features of Kuhn's influential model. In particular, I have focused on what might be called the "prerevolutionary" phenomena of interest to Kuhn, principally (1) the transition from a pre-paradigm state to paradigm-governed normal science, and (2) the nature of normal science itself. Now, one of the worries that Kuhn expressed in the wake of the unbridled enthusiasm surrounding *Structure* was that, due to its "excessive plasticity," *The Structure of Scientific Revolutions* "regretfully . . . can be too nearly all things to all people" (Kuhn 1977, 293). Because Kuhn thought of himself not as a tripped-out 1960s cult leader but rather as a scholar intending to reveal something peculiar about the nature of knowledge and the role of history therein, the wide and unwieldy embrace of his ideas was troubling in that it seemed to indicate that those ideas were so vague that they could easily be instantiated by sheer chance in a vast array of highly disparate kinds of systems. Instead of having gained deep insight into how scientific knowledge develops over time, perhaps in the end all Kuhn had done was coin a few terms whose conditions of application were so loose that their ability to successfully describe some phenomenon really offered no insight at all. In light of this worry, and in light of the fact that the image of science

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

offered here is strongly informed by Kuhn, we must address a serious question about the substantive intellectual content of the proposed analysis—namely, what reason do we have to think that the resemblances between science and adaptive evolution for which I've been arguing over the past few chapters are substantive and worthy of further investigation and development, rather than simply an artifact of comparing side by side a sufficiently wooly description of the evolutionary process with a sufficiently wooly description of the nature of scientific inquiry?

There is a larger question here about the use of history as a way of gaining some kind of more general understanding. While it would be ludicrous to deny that general causal processes can sometimes make history repeat itself, it is equally sinful to attempt to discern a pattern in every set of historical data one encounters. Previous generations had tended toward this latter sin; our time is one in which, regrettably, the search for historical generalizations is very often viewed as naïve. It is true that, as Kuhn worried, patterns can always be found if we sufficiently loosen our criteria. But those who currently view historical generalizations as "not good history" would do well to remember that looking at too fine-grained a level can obscure genuine causal patterns. No trend is visible from a single data point. No forest is discernible in a single tree. Like many trends, a forest is more than a group of things we decided to name. It is a real thing that causally explains lots of natural phenomena. There are many things we will fail to understand if we ignore the forest level or treat it as an epiphenomenon. The two main challenges are (1) to strike the right balance between historical detail and explanatory abstraction, and (2) to be able to distinguish between genuinely explanatory abstractions and those that are vague enough to abuse.

Both challenges can be met simultaneously by first reflecting on what we should expect to see in the historical record if our causal model is correct, and then checking the historical record to see whether those expectations are realized. Looking systematically at an array of independent implications of a model lowers the chances of a coincidental match between the model and reality, while the effort to make our historical explanations empirically accountable tells us how much historical detail we need to adequately support causal historical assertions. Kuhn's quite detailed causal model has been subject to a multitude of such "tests" since its publication. In general, its record of success is mixed, with some implications (like transition to

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

paradigm-governed inquiry and normal science) faring extremely well and others (like incommensurability) less so (Hacking 2012).

What does this challenge look like with respect to our current investigation? Plainly, we need to derive a set of predictions for what we should expect to see in the history of science in contexts where it is governed by the Darwinian evolutionary process. We first took up this challenge by using Godfrey-Smith's concept of a paradigmatically Darwinian population to fix some basic constraints on how scientific inquiry would have to be structured in order to be susceptible to the influence of natural selection. Happily, it turned out that the phenomena that Kuhn identified as causally relevant to the progress of science can often be mapped onto the different dimensions of a Darwinian population in a way that explains why scientific practice in a particular community is given to the degree of directional refinement—vertical progress—it tends to exhibit. I am convinced that this overlap is no accident; these phenomena captured Kuhn's attention precisely because they are what qualifies scientific inquiry as a system susceptible to modification by natural selection. I am also convinced—although somewhat less so—that the fact that Kuhn himself believed the historical development of scientific practice to be governed by natural selection is also not an accident. The three facts—(1) Godfrey-Smith's Darwinian constraints, (2) Kuhn's causal explanation of progress, and (3) Kuhn's belief that the analogy with Darwinian evolution was "nearly perfect"—are independent of one another and yet mutually reinforcing.

The analogy's perfection seemed to increase as we attempted to gain a more precise comparison between the historical development of scientific practice in some instances and the Darwinian process, looking in detail at the way in which the contours of the paradigmatically directional adaptive walk fit those described by the tightly governed march of normal science. In particular, a focus on the developmental level that interested Kuhn—that of truly epic transitions in the history of science—revealed that the relative ordering and magnitude of the independent steps, which together constitute such transitions, are in essential agreement with those of an adaptive walk. We were also able to offer a philosophical explanation for why these phenomena are so similar: in both domains, directional movement is generally not possible until the appearance of a change large enough to define a state space and the constraints that govern movement within it. Once that

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

occurs, subsequent changes are drawn from a new, highly restricted pool of possibilities that undergoes further restriction as new changes become fixed across the population.

We relied on Orr's theoretical model of an adaptive walk which, for compelling reasons, implied that adaptation to a local optimum exhibits certain scale-invariant properties. These properties provided a further opportunity to see whether the evolutionary model of scientific inquiry was credible—specifically, that same historical pattern realized at the scale of Kuhn's truly epic transitions should show up at much smaller scales of change in scientific practice. Through the concept of the *foothold*, we were able to see a variety of very different contexts in which the signature pattern of the adaptive walk is instantiated at relatively small scales. The fact that footholds are crucial in many different contexts is important. Contained within every large-scale transition are myriad small-scale transitions of different kinds. Thus, we should expect the adaptive walk to be reflected in an array of fundamentally different parts of the process of inquiry. And indeed it is.

We thus have a veritable cornucopia of independent lines of qualitative evidence pointing toward the conclusion that the refinement of scientific practice often proceeds under the directional influence of natural selection. I will offer no further arguments on this question. I take the evidence presented thus far to be sufficient to establish that there is often a Darwinian core to the progress of science along the vertical dimension—that is, the directional *refinement* of scientific practice.

Armed with the presupposition of adaptive directionality, I turn now to two higher-level evolutionary phenomena and their instantiations in science. Chapter 5 examines (what I earlier referred to as) directionality along the horizontal dimension—that is, the evolutionary phenomenon of *branching*. Chapter 6 ventures beyond Darwin to take up the problem of mass extinctions. Extending the evolutionary model of scientific inquiry to accommodate the phenomena of branching and mass extinction is an important trial for the model's adequacy. For, after adaptation, branching and mass extinction are the most significant evolutionary phenomena; they leave what are by far the most distinctive imprints on evolutionary history.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

5 Branching and the Growth of Scientific Knowledge

Our primary task in this book has been to investigate the extent to which the evolutionary process might serve as an apt framework for modeling the historical development of scientific knowledge. The larger purpose of this investigation is twofold. First, there is the question of whether and to what degree the evolutionary process of directional refinement is adequate as a description of the causal process that generates scientific knowledge in certain well-described cases. Second, the ability of the evolutionary process to successfully capture the dynamics of scientific knowledge is to serve as a philosophical explanation for why scientific knowledge often exhibits a pattern of development that distinguishes itself from other forms of human activity in a way that is taken to be exemplary of the phenomenon of epistemic progress.

At the core of our efforts lies the concept of a Darwinian population. The degree to which a group satisfies the criteria for being a Darwinian population is to serve as a measure for how effective natural selection is (or could be) at modifying it; or, to put it in slightly different terms, it is a measure of natural selection's particular contribution to the overall modification of a group. Throughout the last few chapters we have used these criteria to make various claims regarding what it would look like for scientific practice to approximate the epistemic equivalent of a paradigmatically Darwinian population—what I have called an idealized epistemic community. We've also looked at aspects of scientific practices in some communities that often seem to closely instantiate the relevant properties of a paradigmatically Darwinian population. Chapter 2 showed that there are several good reasons for thinking that scientific practice often appears to satisfy to a high degree the *heredity* criterion. Chapter 3 likewise argued that there are good

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

reasons to think that scientific practice often appears to satisfy the *intrinsicality* condition. Finally, chapter 4 argued that a population undergoing optimization by directional natural selection will display a characteristic sequence of modifications of related sizes, and that the historical development of scientific practice often appears to fit that characteristic sequence very well indeed.

This last step allowed us to connect the evolutionary model's potential success as a causal explanation for the historical development of scientific knowledge on the one hand with the model's promise as a normative explanation for the special epistemic status of some instances of scientific inquiry on the other. We take science's ability to solve problems as partly constitutive of what makes science epistemically special. And its ability to solve problems in the way that it does is explained by the fact that it possesses versions of the same properties that cause biological populations to solve the problems posed by their environments. It thus turns out that part of our conception of what makes science epistemically special is an inevitable consequence of the causal process responsible for the historical development of scientific knowledge.

We are now faced with a new explanatory task which, like the one previously discharged, contains a descriptive dimension, a normative dimension, and a dimension that seems to be at once descriptive and normative. In general, the success of the evolutionary model as a normative account of scientific knowledge depends on whether we can derive science's epistemically special features as inevitable consequences of the fact that scientific practice, because of its approximations to a paradigmatically Darwinian population, is often highly susceptible to modification through natural selection. We have done this for science's ability to solve problems. Now we must do the same thing for the perpetually increasing breadth of scientific knowledge. That is, we need to show (1) that the modification of scientific practice through natural selection in idealized scientific communities inevitably results in the broadening of scientific knowledge; (2) that the broadening of scientific knowledge, like problem solving, has a direct biological analogue; and (3) that this biological analogue is also the direct and inevitable result of natural selection's modification of paradigmatically Darwinian biological populations. For simplicity's sake, I use the term "growth" throughout this chapter to refer specifically to the increasing breadth of scientific knowledge.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

5.1 The Epistemic Significance of Growth

Some historians of science believe that Isaac Newton probably possessed all of the scientific knowledge that existed during his lifetime (for example, Westfall 1980). As copious as Newton's mind surely was, not even he could achieve mastery over more than a fraction of the body of knowledge to which science can now lay claim—and not just because he is dead (although that is a big part of the problem). Since Newton's time, scientific knowledge has expanded to such an extent that almost all practitioners are well versed only in the literature that pertains directly to the highly esoteric area of research in which they've chosen to work. Although practitioners probably have a general knowledge of all the branches of a field that they need to study before getting to their esoteric corner of research, their knowledge will tend to be fairly basic. They will usually not know much about the frontiers of current research in areas other than their own.

The unwieldy expansion of scientific knowledge since the late seventeenth century is perhaps just one particularly strong illustration of the fundamental fact that knowledge grows over time. But why is the fact that knowledge grows over time epistemologically significant? That is, why should we believe that the mere phenomenon of epistemic growth might be able to tell us something about the nature of knowledge? I think we can articulate at least four independent reasons for its significance, each of which suggests a separate lesson for the study of knowledge itself.

The first reason is simply that any historical trend exhibited by knowledge is prima facie epistemologically significant. The formulation, development, and refinement of a theory involves responding to the factors that function as constraints on a theory in the relevant domain. This is as true of philosophical theories as it is of scientific ones. Philosophical theories of knowledge have, in the past century, taken as their primary set of constraints some intuitions about descriptions of cases in which certain allegedly necessary conditions on knowledge are satisfied (e.g., Gettier cases) (Gutting 2009). Now, regardless of whether or not that particular species of constraint is effective, the intention behind its use is to sort potentially viable philosophical theories of knowledge from the weak, the nonfunctional, and the generally unpromising ones. The viable theories obey the constraints that are recognized by a research community at a given time. Some forms of progress in epistemology (as in other areas of theorizing)

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

involve refining our theories in response to such constraints, uncovering new constraints, and discovering that some constraints had been adopted erroneously.[1]

The epistemological significance of knowledge's historical trends emerges through the ability of these trends to function as constraints on a theory of knowledge. A historical trend offers us a picture of how knowledge behaves (1) on large time scales and (2) at the community level, both of which reflect in their own way the close association between knowledge and some form of *stability* (Kitcher 2015, 477). Trends that hold across relatively large chunks of history reflect the insensitivity of certain aspects of knowledge to perturbations of any kind, because these histories are assumed to include a panoply of different sociocultural, intellectual, and environmental conditions. Because variation in these conditions evidently does not disrupt the relevant epistemic phenomena, we impute to such phenomena an ability to persist despite radical changes of context. We are thus given some sense of what knowledge behaves like regardless of the conditions that prevail at any given time. Whether or not these trends resonate with our intuitions about knowledge, they seem as good a candidate as anything for serving as constraints that a viable theory of knowledge would have to obey. To echo Kuhn (yet again), "how could history of science fail to be a source of phenomena to which theories about knowledge may be legitimately asked to apply?" (Kuhn 1962b, 9).

The second reason for the epistemological significance of the growth of knowledge over time is that growth is often treated as an index of success. Success, in turn, typically demands a causal explanation; we are reluctant to accept accounts of success that are grounded in brute luck. The success of scientific knowledge, manifested in its spectacular growth since the seventeenth century, is no accident; the growth of knowledge is epistemologically significant in that it reflects the ability of knowledge to cause successful outcomes.

To make the point sound somewhat less trivial, consider a counterfactual history in which human knowledge had a tendency to remain more or less static, or perhaps even to shrink. Would the conceptual connection between knowledge and success still hold under these conditions? It is not clear to me that it would. If, for example, knowledge remained static for a relatively large chunk of history regardless of intense investment in rational inquiry, it would then appear to be the case that knowledge is not particularly notable

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

for its contribution to the success of our intellectual endeavors. Thus, despite the intuitive appeal of the idea that knowledge facilitates successful outcomes, this counterfactual history would offer prima facie evidence to the contrary. From this vantage point, the historical fact that knowledge does (thankfully) grow over time is deeply significant. It provides strong justification for the use of the causal connection between knowledge and success as a constraint on philosophical theories of knowledge.

Taking the basic association between growth and success as our starting point, we can distinguish a couple of different forms of success, only for one of which is growth a symptom. Over the last few chapters we have been looking at how ideas might be directionaly refined and perfected in a manner consistent with varying degrees of influence of natural selection. The relevant sense of *success* in that context was one where researchers were able to develop better answers—and better questions—by building on the contributions made by previous researchers to the understanding of a given phenomenon. But success in this sense is perfectly consistent with zero growth. Indeed, under perfectly normal circumstances our total corpus of knowledge might actually contract during—and possibly as a consequence of—the process of refinement. For example, in chapter 3 we discussed how there can often be selective pressure to reduce the number of different kinds of phenomena one is studying because too much variation can disrupt the optimization process. Thus, it is conceivable—even expected—that our knowledge might sometimes cease to grow or even contract as we attempt to deepen our mastery over a particular phenomenon, because we might have either (1) temporarily suspended the investigation of new phenomena (stasis) or (2) abandoned some of our working posits regarding certain phenomena in lieu of a less general but more precise solution to a particular problem (contraction) (Kuhn 1962b, 170). Yet there is a clear and distinct meaning to the notion that we have achieved success in this instance; we had a good solution to a problem, and we've improved on that solution.

By contrast, the sense of *success* for which growth is a symptom needn't indicate any sort of refinement. Rather, success of this kind involves the broadening of our understanding through the investigation of hitherto unexplored or unknown phenomena. In such cases, our success resides not in the development of a completely satisfying account of a new phenomenon but in *the development of a new line of inquiry*. The mere existence of this new line of inquiry provisions future generations of researchers with

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

an opportunity for the kind of refinement-oriented success described in the preceding paragraph. Research communities who achieve success of this kind will often have only a vague impression (if any) of how to approach the object of inquiry. But their initially tenuous grip can (and in many instances does) strengthen over time, fueled by the directional pressures detailed in chapter 4.

As promised, that chapter attempted to draw a connection between researchers' tightening grip over a phenomenon on the one hand and an instantly recognizable sense in which science makes progress on the other. This sense is one of two distinct but overlapping varieties of scientific progress—namely, the movement toward an optimum in practice space that we earlier labeled "vertical progress." If we think of progress in this sense as a kind of movement from one location in practice space to another location, the relevant notion of *success* is that of occupation of such a location. In general, success is something that happens at a specific time, whereas progress is what happens over time. The diachronic trend denoted by "progress" is composed of a series of synchronic events, each denoted by the term "success."

This relation between the diachronic phenomenon of progress and the synchronic phenomenon of success generalizes to the present chapter because the growth-centered species of success is precisely the reason that growth and progress are treated as conceptually linked in some contexts. If the growth-centered species of success is the addition of a new line of inquiry, there is a related growth-centered species of *progress*, constituted by a positive trend in the number of phenomena of which we have at least a dim understanding. This form of progress—which I earlier called "horizontal"—is associated with the persistent tendency of rational inquiry to develop new branches of knowledge, to add "twigs" to those branches, and to add yet further, nubby outgrowths to those "twigs," invoking the growth of the proverbial tree of knowledge over time. Horizontal progress—the overall tendency to increase the number of things we understand or investigate—is the third reason for the epistemological significance of the fact that knowledge grows over time.

Armed with this new species of progress, we can articulate a fourth and final way in which the growth of knowledge bears on a philosophical theory of knowledge itself. We observed above that the growth of scientific knowledge in some cases appears on the whole to be insensitive to contingencies of time and place, for which we awarded epistemic growth the highly desirable property of *stability*. Of course, some forms of stability are

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

more desirable than others; not all forms are epistemologically significant. For example, some of the temporal variation over which the growth of knowledge remains stable involves radical changes in which styles of dress are popular. That's not very interesting or surprising from an epistemological perspective, because our (warranted) default presupposition is that styles of dress are not all that relevant to knowledge.

But there is another sort of change across which the growth of knowledge is historically stable, one with which any philosophical theory of knowledge must reckon—namely, the content of our scientific theories. In Ian Hacking's words, "we know more about polyhedra or atomic weights than we once did, even if future times plunge us into quite new, expanded reconceptualisations of those domains" (1979, 384). Over the past two centuries, we have refined an early insight due to Dalton regarding what elements are and how they combine, but we have also since that time greatly expanded the number of phenomena in this domain that fall within our grasp. This latter trend holds despite radical changes in our views about the nature of atoms and of subatomic particles (e.g., from classical to quantum mechanical). Our ability to increase the number of things we know (or are on the cusp of knowing) has not been impeded by even as drastic a shift as that involved in the transition from discrete to continuous conceptions of matter. Indeed, that ability has almost certainly been enhanced.

What this observation suggests is that one form of progress—horizontal progress—is a higher-order phenomenon that is robust across changes in its constituents. Once we've added a line of inquiry, we have made horizontal progress, even if we eventually abandon everything we initially believed about the phenomena associated with that line of inquiry. Being concerned strictly with the cardinality of the set of phenomena under investigation, horizontal progress generally does not require that our views about its constituents remain fixed. Compare this content independence of horizontal progress with the content dependence of vertical progress. Because the core of vertical progress is the refinement of practice, there must be some causal continuity between earlier and later successes; later successes build on earlier successes. By contrast, progress in the horizontal dimension does not (or, at least, need not) occur by building on previous results. Lines of inquiry can be added quite independently of one another.

To summarize our discussion so far, the historical growth of knowledge is epistemologically significant because (1) historical trends in knowledge

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

function as constraints on a theory of knowledge; (2) growth is causally associated with success, and the success of science is a large part of what makes it appear epistemically special; (3) the accumulation of instances of the kind of success that constitutes growth is associated with a form of progress; and (4) this form of progress is not directly dependent on the content of our practices.

The preceding considerations establish the growth of knowledge as a datum that any credible theory of scientific knowledge needs to imply. Because of its manifest epistemological significance, epistemic growth must function as a constraint on what kinds of philosophical theories of knowledge promise to deliver the most complete understanding of epistemic phenomena. Theories of knowledge that imply epistemic growth as an inevitable consequence of the presence of knowledge are, therefore, better positioned than others to illuminate the nature of knowledge itself.

Our analysis of epistemic growth suggests that we think of it in terms of the closely allied (and conveniently epistemic) notion of *success*, and that the growth-oriented notion of success be understood in terms of the addition of a refineable line of inquiry. This, in turn, led to a conception of *horizontal* progress, which centered around the accumulation of such lines over time (as opposed to vertical progress, which is consistent with the nonproliferation of lines of inquiry). This analysis has raised the bar for successfully linking the growth of scientific knowledge to some biological phenomenon, because it requires that we be more precise about what we mean by "growth" in biological contexts. In particular, the relevant form of biological growth should bear a close resemblance to the sense of success associated with the addition of a refineable line of inquiry.

In section 5.2 I show that the conception of growth operative in evolutionary biology closely mirrors the notion of epistemic growth as analyzed above. This will establish the prima facie case for thinking of growth-oriented epistemic success as an evolutionary phenomenon. I then move on in section 5.3 to consider the evolutionary processes by which lineages undergo lineage addition—what I will henceforth refer to as *branching*. It will turn out that only certain kinds of branching events follow inevitably from the process of natural selection. Since part of what makes idealized scientific communities exemplary from an epistemic point of view is the way in which their distinctive epistemic properties follow from the process of

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

natural selection, we need to be able to model at least some kinds of scientific growth as branching events of these kinds. Following our discussion of evolutionary branching in general, in section 5.4 I examine the arguments for thinking that scientific practice does, at least sometimes, branch in this epistemically desirable manner. In part II of the book, I use an extended case study to illustrate the nature of epistemic branching.

5.2 Evolutionary Success

To establish the claim that the growth-oriented conception of scientific success is a kind of evolutionary success, let us begin by examining the concept of *evolutionary success* as it is (and has historically been) used in paleontology.[2] Warning: it is a bit of a mess.

Two dominant themes emerge from a survey of discussions of evolutionary success in the paleontological literature, both of which can be discerned in the following three excerpts picked at random from the writings of the century's most prominent paleontologists:

> If sheer weight of numbers is the deciding factor, this is the most spectacularly successful of all the phyla, for it includes the insects. (Simpson 1949, 28)

> We know, for example, that some biologic groups have been much more successful than others, as measured by taxonomic or ecologic diversity or by temporal persistence. (Raup et al. 1973, 526)

> This radiation spawned highly successful lineages of woody plants from the standpoint of their [species] longevity, structural diversity, and species numbers. (Donoghue 2005, 85)

One theme is that paleontologists are well aware of the fact that there is no consensus statement on what the term "success" means in evolutionary contexts; most uses of "success" are accompanied by some kind of clarificatory hedge or other. The other theme that emerges is that, despite the known absence of a precise definition, paleontologists seem to more or less agree on what counts as evolutionary success; one rarely encounters disputes over whether the term "success" fits a particular case, or whether certain conceptions of success are incorrect from an evolutionary standpoint. Thus, "success" and its more frequent cognate, "successful," are used fairly liberally, but always with a statement describing the conception of success being invoked.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Paleontologist Pete Ward has recently tried to capture the variegated applications of the evolutionary concept of success in a list composed of eight distinct notions, each of which corresponds to what has been meant by "evolutionary success" at different places in the paleontological literature:

1. *Individual longevity*: how long an individual organism lives
2. *Species longevity*: for a higher-level taxon, the average duration of time between when a species first appears in the fossil record and when it disappears
3. *Species fecundity*: the number of daughter species produced by a given species
4. *Individual abundance*: the number of organisms that are members of a given lineage
5. *Percentage of the planet's biomass*: the proportion of the organic material comprising a lineage, relative to the total amount of organic material on Earth
6. *Species that co-opt other species for their betterment*: species that either manipulate or exterminate other species, resulting in an increase in individual abundance
7. *Wide geographic range*: a lineage for which the members that are farthest apart from one another are far apart in an absolute sense
8. *Surviving mass extinctions*: lineages that were not driven to extinction during at least one of the major extinction events in the history of life

There are a few observations we can make about the list. First is that it doesn't capture every sense of "successful" employed in the passages quoted above. Simpson's "sheer weight of numbers" corresponds to *individual abundance*. Donoghue's "species number" and "longevity," as well as Raup's et al. "diversity or . . . persistence," are all captured by *species fecundity* and *species longevity*, respectively. However, Donoghue's "structural diversity"—that is, the number of different morphotypes in a lineage—cannot be accommodated directly by any conception on Ward's list.

The second thing to notice is that all conceptions (save one)[3] are defined in terms of some form of *relative growth*: relative growth in (1) years (*longevity*); (2) diversity (*lineages* or *structures*); (3) members (*abundance* and *biomass*);

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

(4) area (*geographic range*); or (5) habitats (*geographic range*). Underlying these distinct conceptions is a notion of evolutionary success that seems fundamentally to be about the tendency of a clade to exhibit relative growth along a dimension that is evolutionary significant on long timescales.[4]

Third, each subcategory of success describes a property associated with the *stability of a clade across environmental change*. Longevity is a measure of stability across time, which is itself a proxy for degree of environmental change; the longer the time, the more environmental change there is likely to have been. Abundance is associated with stability across change in a different way, in that it reduces the overall probability of extinction through either sheer weight of numbers[5] or increasing the amount of variation available for responding to environmental change.[6] Diversity and geographic range are both associated with stability across different environments and hence across environmental change. The centrality of stability across environmental change also allows us to capture Ward's inclusion of "surviving mass extinctions" on the list of properties that are closely allied with evolutionary success.

Putting these second and third observations together, it is probable that the reason why these forms of relative growth are evolutionarily significant on large time scales is precisely because they are associated with stability across environmental change. The more growth a clade undergoes in one of these categories, the more likely the clade is to persist across changes in the conditions to which its members are adapted. A stability-centered generalized conception of evolutionary success would then follow: *evolutionary success* involves a clade's relative growth along any dimension associated with stability across environmental change.

We can further improve on this conception by reflecting on a fourth trend—namely, the hierarchical relationship between the different kinds of evolutionary success identified thus far. It will be noticed that each kind of evolutionary success fits into a nested hierarchy of sets of environmental challenges. At the lowest level sits the *abundances*, which correlate to the narrow set of environmental challenges faced by a local population and which, to an extent, define it as a population (see chapter 2). Above that sits *species diversity* and *longevity*, which group populations into distinct lineages within a clade, each of which is associated with a distinct set of environmental challenges that distinguish its evolutionary fate from that of its

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

sister lineages. Moving up a level we find *geographic range*, which includes not only the environmental challenges that a clade's lineages actually encounter, but also those that they could or did encounter—that is, each of the environmental factors that either could potentially affect the selective advantageousness of any trait in the clade or did affect advantageousness at some point as the clade's geographic range expanded. At the very top of the hierarchy rests *surviving a mass extinction*, which includes every possible set of environmental challenges for many effectively unrelated clades.

Each level in the hierarchy corresponds to a qualitative distinction in the degree of difference in the nature of the environmental challenges faced by denizens of a given rank. The problems faced by individual members of a local population will be maximally similar to one another; this is why competition is evolutionarily significant primarily at the population level. Increasing abundance shows that the population is getting better and better at solving the problems that make a difference to survival and reproduction. Next are the sets of problems faced by sister species, which are expected to be different enough to have kept them reproductively distinct, but similar enough to have made branching possible. Following that, the breadth of a geographic range and the time taken to expand it are both metrics for the degree of difference between the problems encountered by members of a clade, with wide ranges indicating potentially radical differences in the kinds of problems affecting members. Finally, those lineages that survive mass extinctions have been exposed to the most radical differences in environmental challenges of which we are currently aware.

Viewed in this way, each kind of success involves the expansion either in degree or in kind of a group's capacity for solving evolutionary problems. Increasing abundance reflects growth in the ability to solve the problems of local adaptation. Increasing diversity reflects growth in a clade's capacity to solve problems that are distinct enough to promote or maintain divergence. Increasing geographic range reflects growth in a clade's capacity to solve qualitatively different kinds of selective problems. And surviving a mass extinction reflects a clade's capacity to rapidly adapt to entirely distinct selective regimes. Incorporating this fourth trend into our generalized conception of evolutionary success suggests the following picture: at any given level, *evolutionary success is distinguished by relative growth in problem-solving capacity*.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

5.3 Assessing the Analogy between Evolutionary Success and Epistemic Success

The formulation of evolutionary success as growth in problem-solving capacity highlights the close connection between it and the growth-oriented sense of epistemic success described in section 5.1 and makes plausible the idea that growth-oriented epistemic success is a type of evolutionary success. Fundamentally, both conceptions of success are centered on the idea that adding members to a group increases the group's capacity to solve problems. In the case of scientific practice, the "group members" are lines of inquiry (or a superset of lines of inquiry; see later in this section). In the case of biological groups, the members are individual organisms, lineages, or groups of lineages.

I think there is cause to take this analogy very seriously. First, the depiction of knowledge as a tree—a collection of "groups subordinate to groups"—dates back to ancient Greece and continues steadily to the present day. These depictions are indistinguishable in form from representations of the tree of life, a tradition that grew out of the effort to classify organisms in a nested hierarchy. In both domains, "growth" has consistently involved the addition of "branches." Second, both biological and epistemic branches are *perpetuated by descent*. In chapter 2 we saw a variety of mechanisms that can work effectively to reproduce scientific practice with a high degree of fidelity. As the emphasis on historical relationships became increasingly important in both domains—histories of science and natural histories—the branching scheme took on a new significance, in that it was intended to represent lines of descent. In the same way that the tree of life grows not by merely adding new lineages but by new lineages emerging from other lineages, so too does the growth of knowledge routinely proceed by the emergence of new lines of inquiry from existing lines. Relatedly, in both cases resemblances between branches are most commonly explained in terms of degree of genetic relatedness. As with species, two branches of knowledge most closely resemble each other whenever they are immediate descendants of a common ancestor.

Third, recall that the growth of knowledge was associated with a recognizable conception of progress in the following way: The addition of a refinable line of scientific inquiry qualifies as an epistemic *success* in that it constitutes a broadening of our understanding through new (or new kinds

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

of) investigations; the accumulation of such instances of success over time constitutes (horizontal) progress—that is, a positive *trend* in the number of (kinds of) phenomena of which we have at least some minimal understanding or on which we now have some degree of traction. How, if at all, does the accumulation of biological lineages correspond to progress in this sense?

In chapter 1 I used some of Darwin's ideas on evolutionary progress to formulate a basic picture of what progress looks like when the historical development of a lineage is governed by natural selection. In practical terms, Darwin imagined a population whose members, despite continuing to improve their ability to profit from the resources for which they were already fairly well suited, had failed to make use of a particular endemic resource—say, a source of potential nutrition—due to the absence of variation capable of exploiting it. If a suitable variant were to arise, the population could expand its resource base, allowing it to grow through its ability to feed more members.

We can, alternatively, envision the expansion of a population's geographic range through similar means: a variant arises capable of exploiting an available resource whose range is broader than the current resource profile for that population. This variant is thus able to venture beyond the population's current geographic range, where it will encounter yet further novel sources of sustenance that it might someday chance to exploit as well. In this way, the original population expands both in abundance and in geographic range. According to Darwin's picture, the descendants of newly emergent variants will—after a sustained period of pursuing hitherto unexploited opportunities—have diverged morphologically from their parent population to such an extent that they eventually comprise a distinct lineage whose evolutionary fate is no longer directly tied to its parent population. He labeled this effect "the Principle of Divergence" and took it to illustrate how natural selection, along with explaining adaptation, is also able to account for the fact that there are so many different kinds of organisms, and that those kinds form "groups subordinate to groups" (Darwin 1859, chap. 4).

Driven by this process of "divergence," the accumulation of biological lineages amounts to a positive trend in the number of environmental problems that a clade is involved in solving. The opportunity to diverge begins with the appearance of a variant that partially solves a problem that has never been solved—namely, how to exploit an available resource. Once the exploitation begins, given enough pressure or time, that variant's descendants will have evolved into their own stable, self-perpetuating lineage.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Through the emergence of a stable lineage in this way, the larger clade of which that new lineage is now part has thus increased the number of problems that it is involved in solving.

At the core of Darwin's image of progress is the phenomenon of *increasing specialization*. Vertical progress is defined by incremental increases in the degree of specialization along a trajectory established by previous generations. Horizontal progress, in contrast, consists of the founding of a new trajectory; where there had previously been no opportunity for specialization, one now exists. Both forms of progress are characterized by improvement in the population's ability to exploit potential resources.

By now it should be clear that the general conception of evolutionary success that was articulated above bears the same relationship to horizontal progress in biological contexts that growth-oriented epistemic success bears to horizontal progress in epistemic contexts. In the same way that scientific communities often appear to progress horizontally by increasing the number of problems that practitioners are involved in solving through the *addition of lines of inquiry*, biological populations make horizontal progress through the addition of new lineages, which increases the number of problems that clades are involved in solving. Fundamentally, the growth of a clade is driven by the emergence of a new opportunity to specialize. And so it is with the growth of knowledge. I believe that there is much to recommend the image of a scientific problem as a resource that scientific communities use to sustain themselves. Problems are the fuel that drives inquiry.

5.4 Evolutionary Branching

This last dimension of the analogy points to the next step in my argument for the idea that part of what makes scientific knowledge epistemically distinctive rests on the fact that the growth of scientific knowledge is driven by natural selection. Up to now, we've been discussing the notions of *growth*, *lineage addition*, and *problem* in a fairly general way, in order to make it easy to highlight the connections among them in both biological and scientific contexts. We now need to introduce a single refinement, common to each notion, that will clarify both their evolutionary and epistemic importance.

Let's begin with the more basic notion of lineage addition. By the addition of a lineage, I mean a *speciation* or *branching event*—that is, the emergence

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

of a distinct, autonomous lineage whose evolutionary fate is not necessarily tied to that of any other lineage.[7] We're not really sure how often this occurs in nature, but it is almost certain that any measurement would underestimate its frequency because most newly emergent autonomous lineages are likely to be of very short duration.

Intuitively, the addition of an ephemeral lineage does not correspond to a kind of growth. This intuition is confirmed by the fact that species longevity is taken to be one of the components of evolutionary success, a notion that we found to be grounded in the concept of growth. We would not count a clade as evolutionarily successful in the case where, although it gave rise to a great many species, each of them quickly went extinct. Rather, in order for the addition of a lineage to count as an evolutionary success—and thus, to contribute to a clade's growth—it needs to last for an appreciable duration on a macroevolutionary timescale.

There is an obvious parallel here with the growth of knowledge. Just as the addition of an ephemeral biological lineage fails to qualify as evolutionary progress, we would not count the addition of a line of inquiry as epistemically progressive unless it made a lasting contribution to human knowledge. When a line of inquiry is judged by a community of practitioners to make a lasting contribution to knowledge, we say that it is *significant* (Haufe ms-b):

> *Def*: a line of inquiry is *significant* just in case it is judged by a community of practitioners to make a lasting contribution to human knowledge.

Now, if horizontal epistemic progress is fundamentally rooted in the growth of knowledge, and if the growth of knowledge is grounded in the addition of a significant line of inquiry, then the accumulation of significant lines of inquiry must be at the heart of horizontal epistemic progress. In other words, one of science's epistemically important attributes is its propensity to generate significant lines of inquiry. Thus, not only is *significance* essential to our understanding of how the growth of knowledge contributes to the epistemic power of science, but the very notion of significance is grounded in the very same property that connects growth to evolutionary success—namely, that in both cases, the addition of a new line constitutes growth only if its effects persist for an appreciable duration.

This is an important result. Yet the centrality of natural selection to our philosophical explanation for the epistemic power of science requires us

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

to go even further. Because (1) if the perpetual emergence of significant lines of inquiry is one of science's most epistemically important attributes, and (2) if the epistemic importance of an attribute is explained in terms of natural selection's causal responsibility for it, then (3) it had better be the case that a line of inquiry's probability of significance—that is, of its making a lasting contribution to scientific knowledge—increases directly with natural selection's influence on the modification of practice.

I want to begin the argument for natural selection's direct effect on the significance of a line of inquiry by looking at analogous phenomena in a biological context, where we already have a relatively firm grasp of the evolutionary dynamics of branching. The essence of the evolutionary branching process is *genetic divergence*. In order for a lineage to undergo an evolutionary branching event—that is, in order for it to add a lineage—the genome of some subgroup within that lineage needs to become different enough to reliably prevent successful breeding with members of the larger parent group of which it is an offshoot.

We can decompose natural selection's role in the branching process into two independent components. First, natural selection might play an important role in branching events by causing certain variants within a population to diverge for ecological reasons. This is Darwin's "Principle of Divergence" discussed in section 5.3, which is known in contemporary evolutionary theory as *ecological speciation*. Second, natural selection might play an important role in branching events when gene flow between two subgroups has been obstructed. I will discuss each of these components in turn.

5.4.1 The Continuum of Ecological Disruption

We can begin to understand natural selection's role in branching events by revisiting the phenomenon of ecological speciation. At the core of the ecological speciation process is, as Darwin clearly articulated, a source of "ecological disruption" (Bolnick and Fitzpatrick 2007). Imagine a population whose sole source of nutrition is grapes. Through directional selection, the ability of members to procure and metabolize Biscuits steadily improves. Meanwhile, a local resource—toast—widely available and rich in potential nutrition, goes unused by the population due to their inability to metabolize it. Suppose that at some point, variants emerge within the population and are able to metabolize toast. Suddenly, there is a decrease in the selective pressure on these variants to improve their ability to procure

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

and metabolize grapes.[8] After all, whatever they lack with respect to the population's average when it comes to grape-eating can be made up with a few easily obtained toast.

At this point, the population begins to undergo what is called *disruptive* or *divergent* selection—two distinct regimes of directional selection, each operating on a different subgroup within the population. While grape-selection continues to directionally improve typical members of the original population, the newly emergent toast-eaters have essentially opted out of grape-selection and are being pushed further and further down the path of toast-eating optimization.

Several factors will affect the chances that an ecological disruption like this will lead to the emergence of an evolutionarily autonomous lineage. The dominant factor is the toast-eating optimum's distance in morphospace. As each generation of toast-eaters continues to improve its toast-eating abilities, the toast-eating subgroup continues to diverge morphologically from its founding fathers and, a fortiori, the grape-eaters. This divergence ceases once the toast-eating optimum has been reached. If the optimum is relatively near in morphospace, toast-eaters will most likely have undergone relatively little modification. And being now only slightly different from grape-eaters, the toast-eating genome will most likely not have diverged enough to persist as a self-contained autonomous lineage. But if (1) the toast-eating optimum is quite far away, and if (2) selection for improved toast-eating is allowed to continue unabated, the chances of permanent divergence begin to look more promising. In short, the further the optimum, the more opportunity selection has to contribute to divergent modification.

Independently of the morphological optimum, another factor affecting the probability of divergence is the expected utility of toast-eating. This is a function of the reproductive benefit of (say) eating one piece of toast combined with the availability of toast. For the first toast-prone variants, divergence will not be possible unless it is at least as beneficial to eat toast as it is to eat grapes. If toast-eating causes a net loss in reproductive success, toast-eating will gradually disappear from the population. But if toast-prone variants are not reproductively penalized for toast-eating, they may exit the grape-eating selection regime and transition to eating toast without negatively impacting their reproductive success.

Since expected utility is a product, we can decompose it into two equally important selective components. One thing that will affect the expected

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

utility is the magnitude of the benefit of eating a piece of toast (not to be underestimated, in my experience). If eating a piece of toast causes an increase in reproductive success, toast-eating variants will become more widespread, other things being equal. But since they are ex hypothesi not competing against grape-eaters, the prevalence of grape-eaters remains unaffected. The toast-eating subgroup is simply allowed to grow due to its ability to exploit an unused resource. Ceteris paribus, the greater the benefit per piece of toast, the faster the toast-eating population grows. The more it grows, the more effective selection will be at optimizing toast-eating morphology, which (as has been mentioned) raises the probability of divergence.

But benefit magnitude can only perform this service if the benefits can be realized. Thus, the other component involves the availability of toast more generally—a proxy for how likely it is that a toast-eater will be able to procure a piece of toast (or several). If there are few toast-eaters and many toast, odds are it's toast for breakfast. If there are many toast-eaters fighting over a single piece of toast, the probability of procuring a piece of toast is very low. Thus, so is the expected utility of toast-eating. And so the probability of divergence varies directly with the probability of toast procurement.

These two components—benefit magnitude and availability—modulate the probability of divergence through their independent effects on how effectively selection can modify toast-eaters. Even if the net gain of eating a piece of toast is relatively low, toast-eating will still be a rewarding strategy if there is ample supply. Equally, if the net gain per piece of toast is relatively high but toast is in short supply, toast-eating will be similarly rewarded. In either case, selection is given the opportunity to optimize toast-eaters on account of the fact that the expected utility of eating toast is at least that of eating grapes.

5.4.2 Gene Flow

The foregoing factors are sufficient to cause a full-scale branching event in sexually reproducing populations under random mating, but only if the value of each factor is on the high end. Otherwise, if grape-eaters and toast-eaters continue to mate despite their divergent lifestyles, the population of toast-eating variants will generally be perpetually pulled back from the toast-eating trajectory and toward the population mean for the ex hypothesi larger parent population of grape-eaters.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

There are a few ways of preventing the swamping of toast-eaters through random mating. Most models of ecological divergence get around this obstacle by positing a nonrandom component to mating, typically in the form of mate preferences (Bolnick and Fitzpatrick 2007, 469). A preferable, more conservative approach[9] would be to introduce nonrandom mating through spacial considerations, along the following lines: toast-eaters like hanging out where there are toast to eat. If toast territory and grape territory share less than 50% of their respective spaces, toast-eaters are more likely to encounter (and mate with) other toast-eaters. Either way, the toast-eating subgroup is able to preserve the progress it has made toward the toast-eating optimum, because its mean phenotype can evolve unencumbered by the effects of the orthogonal optimizing forces governing grape-eaters (Fry 2003).

The simplest and most intuitive model of ecological speciation is, in fact, the process described in section 3.1. Here, assortative mating preferences evolve in response to selection against mating with individuals who might drag one's offspring toward intermediate phenotypic values and thus reduce the fitness of those offspring. It pays toast-eaters to mate with other toast-eaters because that is more likely to advance their offspring toward the toast-eating optimum than mating with grape-eaters. Likewise, it pays grape-eaters to mate with other grape-eaters. In general, "individuals who mate randomly risk producing intermediate offspring with lower fitness, indirectly favoring individuals that mate with their own ecotype" (Bolnick and Fitzpatrick 2007, 470). Those with sufficiently strong mate preferences will be favored by selection, because their offspring will tend to have higher fitness.

Regardless of which model one favors, each of them attempts to devise some way of *interrupting gene flow between subgroups*. It is generally recognized that the cessation of gene flow is a critical element in the development of a variant subgroup into a full-blown new species. Canonical models of speciation are "allopatric," meaning that genetic divergence occurs because subgroups are confined to different areas and mating between subgroups is impossible. Because ecological models of speciation aim at describing conditions under which speciation occurs even when members of divergent subgroups could still mate with each other in principle, such models generally look for ways in which mating between subgroups is permitted but discouraged.

The reason why cessation of gene flow is thought to be so important for speciation has to do with the homogenizing effects of random mating in

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

large populations. This principle is reflected in our discussion above as well as in section 3.1. In such populations, effectively random mating ensures that members' genomes tend to stay relatively similar to one another. Whatever minor genetic deviations arise in a generation will typically be lost due to chance, and so the next generation will more or less resemble its parents.

Sometimes, though, a subgroup becomes geographically isolated from its parent group, cutting off access to the parent's gene pool and the homogenizing effects of breeding within it. Alone in the wilderness, the isolated subgroup's genome begins to diverge from its parent. In such circumstances, there are two distinct gene pools, and the random morphological effects of sampling error play out differently in each pool. In addition, the isolated subgroup—the so-called founder population—is customarily hypothesized to be relatively small. And in a small population, the effects of random genetic drift are more pronounced; a rare mutation that disappears in a large population is less likely to do so in a small population. Thus, in contrast to the ecological process of speciation with which we began, these canonical allopatric models of speciation do not require natural selection to drive populations apart. The amplifying effects of small population size on genetic drift are sufficient to generate permanent divergence between a subgroup and its parent. After enough time has gone by, the isolated subgroup's genome has—purely through drift—become so different from its parent's that successful interbreeding is a low-probability occurrence. A new species is born.

Of course, natural selection also plays a role in fueling the divergence of genetically isolated subgroups. We saw in our discussion of ecological speciation how assortative mating increases the efficacy of natural selection by restricting gene flow between subgroups. Geographic barriers restrict gene flow to similar effect, thus allowing the two subgroups to pursue their distinct evolutionary optima. Although natural selection is not essential for speciation in allopatric models, it is a standard feature of such models because (1) it is ever-present; and (2) its modifying effects are typically much stronger and much more persistent than those of random drift (see section 3.1 for more details).

To summarize, the contribution that natural selection is able to make to the divergent modification of a new lineage increases with an increase in any of the following: (1) distance of the variant optimum, (2) magnitude of the

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

reproductive gain provided by the new resource, (3) relative abundance of the new resource, or (4) restriction of gene flow between subgroups.

As we have seen, another major contributor to the probability of permanent divergence is the small size of founder populations. But unlike (1)–(4), small population size does not increase the efficacy of natural selection; quite the reverse. Thus, the standard picture (right or wrong) of speciation is one in which gene flow is curtailed, thus allowing drift to dominate early on. As the isolated subgroup grows, the probability of sampling error dwindles commensurately, allowing ecological factors to properly dominate the modification of the lineage.

5.5 The Emergence of Specializations

I see the relation that *significant lines of inquiry* bear to *specialization* as that of *well-marked varieties* to *species* in the *Origin*. In general, to say that the emergence of a significant line of inquiry is analogous to the process of ecological speciation in biological contexts is to say that the kinds of Darwinian, selection-promoting factors that generate new species also work to generate specializations of lasting significance to scientific knowledge. What I would like to do in the remainder of this chapter is to provide a schematic description of a selection-driven branching process in science, a process in which the propensity to generate new specializations is modulated by the degree to which the development of scientific practice is dominated by selection—that is, the degree to which scientific practice approximates a paradigmatically Darwinian population.

The aim of this section is to illustrate a phenomenon of central epistemic importance, which often appears to arise in the historical development of scientific knowledge—namely, the inevitable emergence of *significant* lines of inquiry as a by-product of the influence of selection on the modification scientific practice. The emphasis on "significant" is essential here. I do not maintain that every line of inquiry's existence can be attributed to the influence of natural selection in scientific practice. Scientific knowledge does not necessarily benefit from the proliferation of lines of inquiry; it is not a weeping willow, which thrives on the unchecked multiplication of branches per se. Its likeness is rather that of a carefully curated bonsai tree. It must be constantly pruned, staked, and encouraged in very specific ways so that it might have the best chance of satisfying our purposes.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Nor I am suggesting that significant new lines of inquiry can only arise in science through the influence of the sorts of factors that facilitate the efficacy of natural selection. Just as no one working today would argue that new biological species can only arise through ecological speciation, we should embrace a plurality of circumstances out of which significant lines of inquiry have emerged. For example, I acknowledge the presence of serendipity in science, although I think its centrality to discovery has been vastly overplayed. More to the point, I fail to see how anything of epistemic importance might accrue to science by virtue of the fact that it is susceptible to the chance occurrence of significant discoveries. If these discoveries are indeed as chancy as they are often portrayed, why is modern scientific inquiry so disproportionately blessed with them? Why, for example, was the much longer empirical tradition of alchemical investigation not similarly favored by chance?

Ultimately, we must come to grips with and account for the fact that the growth of modern scientific knowledge through the proliferation of significant lines of inquiry is, for some reason, an inevitable consequence of the nature of modern scientific inquiry. This fact, along with the problem-solving capacity of modern science, are its most distinctive epistemic characteristics. The reason for this inevitable proliferation, I argue, is that scientific practice is—at least at times—a nearly paradigmatic Darwinian population and is thus highly susceptible to the influence of natural selection. And one of natural selection's inevitable by-products is the emergence of independent, refinable lineages. We call them "species" in biology. In science, they are called "specializations," or "disciplines."

I am going to argue that a line of inquiry's probability of significance—that is, of its making a lasting contribution to scientific knowledge—increases directly with natural selection's influence on the modification of scientific practice. The case for this claim develops along the following lines: first, I argue for the definition of "significance" stated in section 5.4—in particular, for the part of the definition that makes practitioners' esteem essential to a line of inquiry's significance. I then argue that the sorts of factors that often increase the probability of practitioners' esteem are closely analogous to those that increase the probability of selectively driven (ecological) speciation. What this will show is that the same properties that affect natural selection's influence over the emergence and endurance of new species also affect natural selection's influence over the emergence of significant

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

lines of inquiry. Were we to show this, we will have established that the probability of a significant line of inquiry emerging is directly tied to how susceptible scientific practice is to the influence of natural selection—that is, how closely it approximates a paradigmatically Darwinian population. This supports our contention that natural selection's governance of the historical development of scientific practice is what accounts for its most epistemically significant attributes.

5.5.1 Achieving Significance in Science

We've already offered some discussion of the topic of *significance* in chapter 3. There we made two points that are relevant to the current section. First, we observed that a scientific community sometimes functions as a constructed niche, where the addition of new practices—be they new problems or new problem-solving approaches—is carefully controlled by the preferences of the scientific community. This suggests that, at a very basic level, a line of inquiry is significant when the community decides that it is significant. Indeed, "what better criterion than the decision of the scientific group could there be?" (Kuhn 1962b, 170). Because of the many possible directions in which inquiry could potentially proceed, the decision is not taken lightly. Where a community decides that a line of inquiry is worth pursuing, it has invested in the belief that its pursuit will result in a lasting contribution to scientific knowledge.

We also saw in chapter 3 how science funding agencies attempt to gauge the potential significance of proposed research by examining the strength, proximity, and number of connections it possesses with ongoing scientific research. Although this gauge does not necessarily tell us what significance is per se, it does have two important consequences. First, it provides a constraint on a viable theory of significance—namely, that any philosophical theory of significance should imply that significance is something that can be reliably measured by examining a line of inquiry's connections with ongoing research. Ideally, the reliable detection of significance through an examination of connections with ongoing research should be an obvious, natural consequence of a philosophical account of significance in science.

Second—and related to this constraint—it fits precisely with the Kuhnian notion that significance is an internal matter. Although a community employs some objective criteria when it comes to gauging significance,

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

those criteria are themselves ultimately rooted in the community's prior decisions about which lines of inquiry count as significant—that is, those lines that the community had previously decided would make a lasting contribution to scientific knowledge. This is entirely appropriate from a Kuhnian perspective. Because if (1) the community is better positioned than any other party to judge current significance, and if (2) the relation to current significance is the most reliable guide to future significance, then (3) the community is best positioned to judge future significance. But why, from the community's perspective, should a line of inquiry's probability of making a lasting contribution increase with its relatedness to lines of inquiry that have made, or are poised to make, lasting contributions already?

I think there are a variety of ways in which one could account for the perceived relationship between current and future significance. In keeping with earlier discussions in the book, I will try to explicate this phenomenon within Kuhn's general framework. One of the most powerful consequences that follows from the adoption of a paradigm is that practitioners are furnished with a guide for discerning which kinds of problems are important, which kinds of phenomena are important to understand, which facts will be "particularly revealing of the nature of things" (Kuhn 1962b , 25). In many cases, the paradigm will immediately implicate certain quantities directly, as Darwinian theory did with respect to the mechanisms of heredity, or as Newtonian theory did with the gravitational constant. In both cases, adopting the paradigm consisted partly of accepting that these are questions whose answers are critical for advancing our understanding of nature.

But it is part of the character of scientific inquiry that, as time goes on, new questions emerge out of the investigation of those whose significance was directly highlighted by the adoption of the paradigm. In the case of Darwinian theory, this includes the many questions now pursued by molecular biology—questions that, while directly relevant to the issue of heredity, could not possibly have been conceived of in the mid-nineteenth century. Here, practitioners converged on the significance of these questions because of their direct connection to questions whose significance had been previously implicated. Whether the question's significance is partly constitutive of a paradigm's adoption, or whether it is instead a downstream consequence of a paradigm's adoption, the fundamental

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

factor influencing perceptions of significance is how closely the focal question can be linked to issues that are significant according to the paradigm.

What this reasoning suggests is that to be significant *just is* to be significant according to the paradigm. Significance is, in Carnap's terms, an "internal question"; it makes no sense for members of a scientific community to ask whether the governing paradigm has correctly identified some phenomenon as significant (Carnap 1950). If what allows a group of practitioners to constitute a scientific community is the acceptance of a shared paradigm, and if accepting a paradigm implies accepting certain questions as significant, then one cannot both be a community member and challenge the significance of a paradigm-validated question. In other words, community members must, virtually by definition, accept the significance of paradigm-validated questions.

This, then, would explain both why there can be no higher standard than the assent of the relevant community and why communities judge potential for significance on the basis of relation to past significance. If being significant just means being paradigm-validated, and if paradigm-validated significance only becomes paradigm-validated through the assent of the community, then a question can only achieve significance if the community is convinced that it is significant according to the paradigm. To say that there is no higher standard than the assent of the relevant community is just another way of saying that there is no higher standard than the paradigm. To the extent that we conceive of practitioners as epistemic niche constructors, things could hardly be otherwise.

There is no prescribed route by which practitioners must become convinced of a problem's significance. Its significance may be immediately obvious to all, given its centrality to or ubiquity within paradigm-driven research. Or practitioners may gradually, independently, and quite on their own come to see the relevance of the problem to their ongoing research, whose significance (of course) can be derived from its relation to paradigm-validated questions. Or, as I will explore at length in part II, practitioners may come to be convinced of a problem's significance through explicit argument aimed at establishing that its significance follows by implication from the paradigm. In the final analysis, all routes seem to result in the same basic effect: the community accepts the problem's significance, and it thereby becomes a question worthy of pursuit.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

5.5.2 How Significant Problems Give Rise to New Specializations

Once a community of practitioners recognizes the question as a potential object of legitimate scientific inquiry—that is, as a paradigm-validated question—research on that question becomes part of the broader communal effort to understand nature. At this stage, the community is ready to add a *line of inquiry*. But it is not clear yet whether an entirely new *specialization* will emerge. Whether the new object of inquiry will give rise to a new specialization depends on the same sorts of factors that determined whether a source of potential nutrition would ultimately give rise to a new species.

The disciplinary branching process described below can be understood as a special case of Pagel and Mace's (2004) speciation-inspired model of the production of human cultural diversity. In that model, branching is driven by two primary influences:

> One is the drive to secede from larger groups whenever possible, the better to control some defensible resource; this is what gives rise to the geographical patterns of diversity. The second set of forces is social and behavioural. They maintain cooperation within groups and create cultural identity and coherence, causing barriers to gene flow and meaning that vertical cultural transmission dominates. (Pagel and Mace 2004, 278)

In science, the "defensible resource" is a certain supply of scientific problems. The "social and behavioural [forces that] create cultural identity and coherence" are the replication and isolation mechanisms that allow scientific communities to persist over time. In the reminder of this chapter, I show how these two forces conspire to produce a version of the Pagel-Mace cultural branching process in science.

Whose job is it to solve the new and important problem that stands before the community? We can think about this question in the context of the grape- and toast-eaters we met earlier in the chapter. Whose job is it to eat the toast? Presumably everyone would *like* to be able to eat the toast. But in that population, only certain variants are equipped with the capacity to metabolize toast, and so the toast bonanza is left for them alone.

Toast vs. grape is an extreme case; the canonical Darwinian process is not one in which qualitatively new abilities arise in a single generation. But it is not an impossibility, and it illustrates an important evolutionary lesson—namely, branching events are most likely to occur when there are no intermediate phenotypes (Bolnick and Fitzpatrick 2007, 469). When

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

toast-eating ability is something that one either has or doesn't have, it becomes possible for two selection regimes to emerge, one pushing some variants toward a toast-eating optimum, the other pushing some variants toward a grape-eating optimum. In contrast, if each member of the population can eat both toast and grapes, there will be just one selective process (at least, with respect to toast- and grape-eating)—the selective process that drives the population to the optimum toast-grape balance (Bürger 2005).

The same lesson applies in the context of a scientific community's new research problem. Even if the new problem is important enough to warrant a new line of inquiry, it might not be distinct enough to constitute a new kind of research problem; it may simply be a hitherto unencountered version of the kind of problem practitioners are adept at solving. Indeed, enormous amounts of intellectual energy are often spent trying to transform queer new problems into the kind of problem for which we have well-worn methods of solving. If the new problem is something that could be solved with the community's existing set of practices, or with a variation on existing practices, then the emergence of a new specialization is improbable. When an existing problem-solving method (or some variation thereon) can be applied with success to the new problem, no practitioner can claim privileged access to the methods required to solve the problem. Anyone who can add even numbers together can just as easily add odd numbers together.

We can put the point in evolutionary terms: If the methods of solving the new problem are close enough in practice space to reach via a sufficiently plastic application of existing methods, there won't be the sort of disruptive pressure on the community that is required to split it into ecologically distinct subgroups. But if qualitatively distinct problem-solving methods are required, an autonomous new specialization may be able to emerge. For, at that point, each practitioner will have to choose to acquire one set of methods or another.

Now, because disruptive pressure emerges as a result of the unreachability of certain points in practice space, and because problems are themselves an element of practice, it ought to be the case that sufficiently distinct research problems are capable of driving the emergence of new specializations all by themselves. That is, even if a new problem is solvable with existing methods, it still might be capable of fueling the emergence of a new specialization. This phenomenon is easiest to illustrate in the context of very general problem-solving methods like mathematics and statistics.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Two statisticians who receive exactly the same training may nevertheless find themselves in two different specializations; one might be, say, a data scientist, while the other is a biostatistician. Despite drawing on the same problem-solving toolkit, the problems involved are themselves distinct enough to be unrelated. Progress in data science could easily have no bearing on progress in biostatistics, because there might be very little overlap in their respective domains. This would be a version of the model of ecological divergence discussed above that uses spatial considerations to interrupt gene flow: a population can branch into two distinct species, so long as there are two subgroups whose territorial overlap is sufficiently small.

Mulkay (1975) observed that the introduction of new technologies can have a similar disruptive effect. Given an instrument whose competent operation is sufficiently recondite, it can sometimes happen that entire specializations grow around the use of that technology per se. This is becoming increasingly common, as the variety of tools necessary for answering certain kinds of questions becomes increasingly disparate. Under these circumstances, it will often be impracticable to have a single practitioner develop competence in the use of a certain technology in addition to mastering the subject matter to which that technology is to be applied.

In sum, as with biological populations, the ability of natural selection to drive the emergence of a new autonomous scientific specialization is directly related to how specialized one needs to be in order to extract a benefit from the hitherto unexplored domain. Whether it is a new problem, a new technology, or a new theoretical framework, the less common ground there is between practitioners who engage in one practice and practitioners who engage in the other practice, the more probable it is that the single community of practitioners will fissure into two specializations.

In section 5.4.1 we saw how the magnitude of the reproductive gain provided by toast would have to exceed that provided by grapes in order to lure toast-prone variants away from grape-eating. When the benefit of toast is less than that of grapes, toast-eating is selected against; if it is equal to grapes, toast-prone variants will be indifferent between toast and grapes, and the probability of the lineage splitting is minimal. As the benefit of eating toast increases, natural selection will increase its bias in favor of toast-eating, and the prospects for permanent ecological disruption begin to rise.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

What should we expect in science in light of this dynamic? Let us begin with the observation that, ceteris paribus, the probability of ecological disruption increases as the expected utility of improving one's grape-eating ability declines. As grape-eating optimization becomes less beneficial, the relative benefit of eating toast will climb, causing increased selection for toast-eating. This suggests the following principle for the emergence of new scientific specializations: *the probability that a new specialization will emerge varies inversely with the marginal utility of progress in an established specialization*.

Earlier efforts to analyze the growth of scientific knowledge describe the emergence of new specializations in precisely these terms. Mulkay's (1975) account, for example, begins with the following observation:

> The growth of a new area starts with the perception, by scientists already at work in one or more existing areas, of unsolved problems, unexpected observations or unusual technical developments, the pursuit of which lies outside their current field. Thus the exploration of a new area is often set in motion of a process of scientific migration. . . . Scientific migration is not a random process, for the scientists moving into a new field tend to come from other areas with specifiable characteristics. In particular, they come from research areas which have experienced a pronounced decline in the significance of current results; from areas where there are few or no avenues of research easily available; from areas whose members have special competence in or knowledge of techniques which appear to have wider application; and from areas which have been disrupted, often by events originating outside the research community, and whose members have consequently no firm commitment to an established field. (Mulkay 1975, 5)

A similar dynamic is depicted in the graph produced by Holton (1962) (figure 5.1). Holton analyzed the emergence of new specializations in terms of the change in the "number of undiscovered interesting ideas" as a function of time. Although this is not, strictly speaking, the same thing as "marginal utility of progress," both notions attempt to capture the same idea: the probability of making major progress in a given specialization declines with time. It will be noted that, in his representation, branching events occur right at the point where the slope of the "parent" curve starts to level off—that is, at the point where the probability of making major progress begins to decline. As we will see in part II's case study, the branching point for paleobiology as a new specialization occurs at the precise moment when the magnitude of contributions to evolutionary theory starts to plateau.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

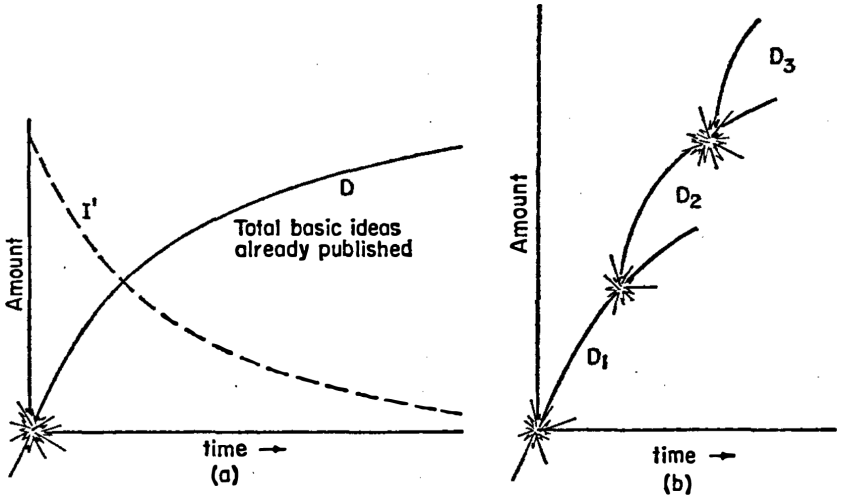

Figure 5.1
The escalation of discovery lines. From Holton (1962).

Both Holton and Mulkay discuss the emergence of specializations as a branching process, one that explains how scientific knowledge grows over time. Mulkay, for example, concludes his account by stating, "This account of the processes of scientific development emphasises the way in which science grows through the *branching of new lines of research*" (Mulkay 1975: 5–7; emphasis added). Holton (who literally uses a drawing of a large, leafy tree to illustrate the process) concludes that "the growth of scientific research proceeds by the *escalation* of knowledge—or perhaps rather new areas of ignorance" (Holton 1962, 126). For Holton, "escalation" occurs through "the applicability of the early techniques, and the originating of new questions now suggesting themselves in neighboring parts of the same fields, [which] provoked a rapid branching into several new directions" (118).

The key point for our argument lies in the fact that, according to these models (and soon to be corroborated by part II), new scientific specializations often appear to arise at the exact moment at which natural selection becomes most effective—that is, when the marginal utility of progress in an established discipline goes into decline. Once the magnitude of the benefit of making headway in a given specialization starts to level off, practitioners who "have special competence in or knowledge of techniques which

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

appear to have wider application" begin to look for opportunities to opt out of the race to squeeze the last few morsels of progress out of a declining specialization. These opportunities can be found most easily in research areas that offer benefits comparable to those available in their current specialization but have yet to see an influx of capable researchers.

This process, in turn, raises the related issue of whether the hitherto unexploited source of problems is sufficiently abundant to support a group of practitioners large enough to propagate itself. On the face of it, this problem is just that of toast-abundance, instantiated in a scientific context: it doesn't matter how important the new type of problem is if there are not enough tokens to allow the community to grow. But scientific problems can behave somewhat differently in the sense that working on them will often generate more problems, such that the growth of knowledge can often really properly be characterized as a steady decrease in the ratio of solved problems to unsolved problems. This is what Holton alludes to when he describes the growth of scientific research "by the *escalation* of . . . new areas of ignorance."

The upshot of this important difference in how the supply of problems behaves is that problem supply will rarely affect natural selection's power to modify scientific practice to the same extent that nutrient or mate supply will affect selection's power to modify biological populations. Even if opportunities in another domain are initially limited, their pursuit can often result in an increase in the supply of problems, thus laying the foundation for the migration of a far greater number of researchers, in Kuhn's words, "away from competing modes of scientific activity." Thus, at least with respect to resource abundance, the influence of natural selection can be counted on to be less constrained in scientific contexts than in biological ones.

The scientific community's ability to restrict the inflow of new practices was discussed extensively in chapter 3, as a way of arguing that new practices often become widespread in science chiefly by virtue of the goodness of fit between those practices on the one hand and the epistemic challenges that scientific communities set for themselves on the other. Thus, much of what I will have to say about the flow of variation per se has already been argued. Here I will simply highlight a few points about the isolation of communities specifically as it pertains to the emergence of new specializations.

As Kuhn made clear at various points, the isolation of a scientific community was the most important factor for promoting its ability to specialize. By definition, specialization requires limiting the scope of inquiry and

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

the tools employed therein. The normative constraints that paradigm adoption places on acceptable problems and methods achieve this isolation. But Kuhn's view that isolation was caused primarily by lexical divergence was overly narrow (see chapter 3 for a critique). If Kuhn's view were correct, there should be far fewer specializations than there in fact are, since it is false that the only way to effectively isolate practitioners is through communication breakdowns. But this is an entirely separate matter from that of whether specialization accounts for the directional refinement of scientific practice, which it arguably does.

Because, *pace* Kuhn, scientific communities have so many mechanisms for isolating themselves from the influx of migratory practices, they can often be rather ideal systems for generating new specializations. Unlike biological populations, which are usually at the mercy of insurmountable physical obstacles to prevent gene flow, scientific communities can police the flow of practices by arbitrary fiat. This gives each community virtually unlimited power to refine methods, problems, and the like, solely in response to the epistemic challenges they choose to confront. The freedom to pursue their own epistemic ends in relative isolation accelerates the divergence of subgroups of practitioners from one another, eventually resulting in the creation of independent, autonomous research communities. I should emphasize that Kuhn largely embraced this image of the branching process, the one difference being that he seemed to think isolation was more difficult than it probably is.

The goals of this section have been (1) to present an interpretation of what it would look like for natural selection to drive the emergence of new scientific specializations and (2) to argue that selection's actual influence on the emergence of new specializations is probably significant in some cases. The argument for (2) is based on two premises, the first of which was our interpretation of adaptive branching in scientific contexts. This interpretation consists of drawing a direct connection between the variety of conditions that independently contribute to selection's influence over the branching process in biological populations on the one hand and the factors thought to be of the most general causal relevance to the emergence of new scientific specializations on the other. The second premise is simply that this connection is no accident. In other words, the reason why factors that often cause the emergence of new scientific specializations are directly analogous to

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

factors that amplify natural selection's role in the emergence of new species is because the emergence of new scientific specializations is often driven by natural selection. Chapter 4 emphasized how the process of refinement of scientific practices is analogous in character to the process of refinement of structures by natural selection. We now see that effect of selection on the growth of scientific knowledge is equally important. When selection's role in scientific inquiry is less pronounced—that is, when scientific practice is somewhat less than a paradigmatically Darwinian population—the practitioners' practices will tend not to diverge from one another with the pace, magnitude, and shape that they characteristically exhibit.

5.6 Conclusion

Under the overwhelming influence of natural selection, the emergence of new scientific specializations is as inevitable as the emergence of new species. Now, if the influence of natural selection on the development of practice makes new specializations inevitable, and if scientific practice becomes more susceptible to selection's influence as its approximation of a paradigmatically Darwinian population increases, then the fact that scientific practice is a nearly paradigmatic Darwinian population would be a very good explanation for its unrelenting propensity to generate new specializations.

We are then forced to take even more seriously the philosophical claim that part of what gives modern scientific inquiry its epistemic power is the fact that its development over time seems often to be dominated by the influence of natural selection. As with the directional refinement of practices, the growth of scientific knowledge is a datum with which any philosophical account of scientific knowledge must reckon. If, as I have argued, the perpetual growth of scientific knowledge consists of the perpetual emergence of new scientific specializations, then the growth of scientific knowledge is inevitable if scientific practice resembles a paradigmatically Darwinian population. There is no available alternative theory that makes science's two most significant epistemic features an inevitable consequence of the process of scientific inquiry under a few specifiable conditions.

And yet there is a further story to tell. So far, our discussion has remained very close to two traditions of thought. At its most general level, our goal has been to shed philosophical light on the historical pattern of development in some instances of scientific practice, and we have relied largely

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

on Thomas Kuhn's descriptions of the various elements of that historical pattern. There have been a few instances in which it seemed necessary to amend Kuhn's views to a degree, but none of those emendations called into question a broadly Kuhnian perspective on what the historical pattern of scientific development actually looks like in certain contexts. Parallel to this discussion have, of course, been our repeated appeals to the tradition of thought known as Darwinian evolutionary theory, whose wide exploration was made possible by our attempt to follow the implications of thinking about scientific practice as a Darwinian population.

These two traditions have, up to now, lined up very nicely indeed—much more tidily than I would have ever dreamed when I first began thinking about these issues. But there are major features of each tradition that we have thus far neglected. It will no doubt have occurred to the reader that, even though Kuhn's model is specifically intended to describe the process of revolutionary scientific change, I have made virtually no mention of any such change, nor have I attempted to find an evolutionary analogue to scientific revolutions.

Does my silence on this matter stem from skepticism about the existence of Kuhnian revolutions? No. I am mostly convinced that the kind of revolutionary change described by Kuhn—incommensurability and all—actually exists. Is it because scientific revolutions lack epistemic significance? Again, no. They are as much a part of the historical pattern of scientific development as either directional refinement or branching. More importantly, major revolutions are viewed by practitioners as epistemically unique, in that there are no other occasions in which so much scientific progress is made in so little time.

However, in contrast to the other epistemically significant phenomena we've analyzed, major scientific revolutions do not follow inevitably from the extent to which scientific practice is a paradigmatically Darwinian population. That is, a population of practices that is maximally susceptible to modification by natural selection does not thereby increase its chances of undergoing a major scientific revolution. All that being a paradigmatically Darwinian population guarantees is that specializations will emerge and that they will be directionally refined through a process Darwin labeled *adaptation* and Kuhn labeled *normal science*. And although Kuhn clearly thought that normal science was a necessary condition for a scientific revolution (more on that in chapter 6), major revolutions require more than

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

normal science. And that extra stuff does not follow as a consequence of paradigmatic Darwinian populationhood.

Which brings us to the feature of evolutionary theory that we've thus far neglected: extinction—mass extinction, in particular. Some of the most significant developments in empirical and theoretical work on evolutionary change since the publication of the *Origin* have been the very recent growth in our understanding of the nature of mass extinction and its fundamental role in shaping evolutionary history. But the subject has yet to come up in this book because, like scientific revolutions, mass extinctions are not an inevitable consequence of the Darwinian evolutionary process. A paradigmatically Darwinian population is no more or less susceptible to a mass extinction event than a paradigmatically un-Darwinian population is.

In chapter 6 I argue that these two phenomena—major scientific revolutions and mass extinctions—are two names for the same thing: a major scientific revolution is a mass extinction of scientific practices. Hence, their simultaneous neglect has not been a coincidence. We will then have to confront the epistemological question of whether major scientific revolutions can be part of what accounts for the distinctive power of modern science, given that they do not follow from the fact that scientific practice is a Darwinian population.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

6 The Recovery of Scientific Knowledge

Up to now we've been examining different components of a cycle that characterizes both adaptation-driven speciation and incremental scientific progress within a discipline. I now want to shift away from this cycle to look at a type of event that, while not part of any cycle (that we know of),[1] is a cycle unto itself, characterized by a predictable series of stages beginning with a catastrophic perturbation and ending with a recovery. The general pattern of these perturbation cycles begins with an abrupt, high-magnitude shift in the set of challenges that a population has been designed to solve. The abruptness of the shift results in the widespread failure of existing solutions. The solutions that emerge during recovery are aimed at solving a very different set of challenges, and consequently look very different from those of the previous regime.[2] In 1962, Thomas Kuhn called these "crises" that usher in "scientific revolutions." In 1963, paleontologist Norman Newell called them "crises in the history of life" that usher in "revolutions in the history of life."[3]

We can think of the process described in chapter 5 as one end of a continuum of branching events, defined by the decreasing influence of selection. Selection will be maximally involved in branching events in cases where branching is caused by selection for divergence (figure 6.1, far left). It will be minimally involved in events where branching is caused by pure drift (far right). We can then map onto this continuum a series of extinction events increasing in size from left to right. This mapping will produce the following rule: *The influence of selection on branching varies inversely with the size of the extinction event*. The present chapter shows how this simple rule can be used to explain a lot about scientific revolutions and about why Kuhn's work on this topic often seems muddled.

As in previous chapters, I will be using a range of correspondences between science and evolution to support the view that Kuhn's major

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

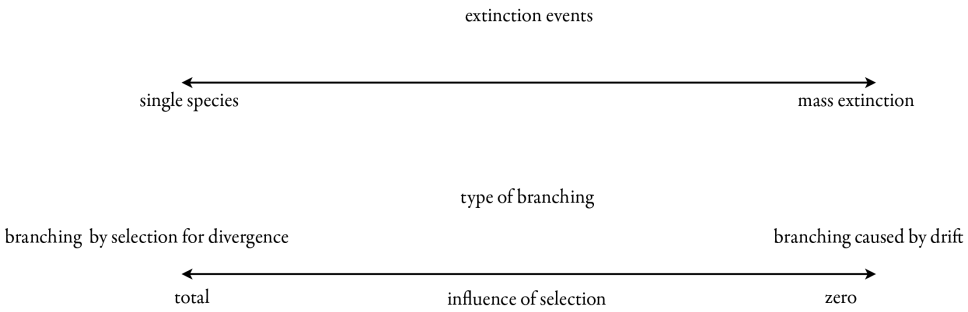

Figure 6.1
The relationship between branching and extinction. Adapted from Brandon (2005).

scientific revolutions are instances of the evolutionary phenomenon known as a *mass extinction*. But our job in this chapter is made considerably more complicated by a number of differences with earlier discussions in the book. First is the fact that Kuhn's views about scientific revolutions remain very controversial, underwent significant revision after the publication of *Structure*, and never achieved an enviable level of clarity. I will have something to say about each of these as the discussion progresses.

The second complicating factor has to do with the non-Darwinian style of evolutionary change that characterizes the mass extinction cycle. We have invested a lot of time thus far exploring the notion of scientific practice as a paradigmatically Darwinian population and showing what sorts of phenomena ought to follow from that status. Because we have staked the epistemic credentials of science on its susceptibility to selection, however, the non-Darwinian nature of mass extinctions—and, by implication, major scientific revolutions—raises the following uncomfortable question: According to my argument, doesn't the fact that the development of scientific practice often exhibits non-Darwinian behavior undermine the epistemic power of science in those cases? My eventual negative answer to that question will end up significantly strengthening the argument that science becomes more progressive the closer it comes to functioning as a paradigmatically Darwinian population.

I begin in section 6.1 by showing how the four essential stages of this evolutionary event are realized in both biological and scientific contexts. Each stage, in turn, has a variety of signatures that show up both in mass

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

extinctions and in scientific revolutions. The resulting set of parallels is truly uncanny. It is worth noting in this connection that mass extinctions did not receive widespread recognition as an important—or even extant—evolutionary phenomenon until the 1980s. Furthermore, the evolutionary theory of mass extinctions from which I will draw in this chapter did not begin to come together until even more recently, much of it after Kuhn's death. And yet the essential components of the mass extinction process as we now understand it are already clearly articulated in Kuhn's 1962 monograph.

In section 6.2 I focus on the difference between the evolutionary process of mass extinction and the more paradigmatically Darwinian processes we've been discussing over the last several chapters. This will lead us in section 6.3 to a philosophical analysis of the epistemic implications of the distinctively non-Darwinian character of some instances of revolutionary scientific change. It will turn out that Kuhn's perspective on the epistemic aspects of scientific revolutions closely parallels the interpretation suggested by an evolutionary model that anchors the epistemic power of science in its degree of susceptibility to Darwinian change. Kuhn made a clear distinction between the epistemic status of changes to practice that occur as a result of normal science on the one hand and those that occur during revolutionary change on the other. Incidentally, this same distinction flows directly out of the evolutionary model. Section 6.4 uses the gains in clarity developed in previous sections to shed some light on certain features of Kuhn's views on scientific change that critics have singled out as particularly bothersome.

6.1 Crisis in the History of Life

> When I had satisfied myself that no star of that kind had ever shone for the before, I was led into such perplexity by the incredibility of the thing that I was not ashamed to admit to doubting the faithfulness of my own eyes; and so, turning to the servants who were accompanying me I asked them whether they too could see a certain brilliant star when I pointed out the place directly overhead. They immediately replied with one voice that they certainly saw it and that it was extremely bright. But in spite of their affirmation, still being doubtful on account of the novelty of the thing, I inquired of some rustic country people who by chance were traveling past in carts, whether they could see certain star on high [*in sublimi*: RW]. Indeed, these people shouted out that they saw that huge

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> star, which had never been noticed so high up. And at length, having confirmed that my vision was not deceiving me, but in fact that an unusual star existed there, and astonished at such a new phenomenon beyond anything that had ever occurred in the heavens, compared with the other stars, I undertook immediately to measure with my instrument.
>
> —Tycho Brahe's account of the *stella nova* of 1572 (quoted in Westman 2011, 233)

> Du fuh?
>
> —(Probably) A particularly observant *T. rex*'s comment just prior to the bolide impact that killed him some 65 million years ago

The *stella nova* of 1572 and the Chicxulub bolide that appeared at the end of the Cretaceous period 65 million years ago are two instances of the kind of phenomenon that characterizes the first stage of a revolutionary transition. Each of these events constituted a "crisis" moment for an incumbent order, one from which it would never recover.

This section will argue through a range of comparisons that major scientific revolutions and mass extinctions are two species of the same genus—a process distinguished by its distinctive cycle of perturbation and recovery. For clarity's sake, I've broken this process down into what I take to be its three major components: (1) an environmental perturbation that causes the failure of incumbent problem-solving capacities; (2) the replacement of incumbents with novel forms of specialization; and (3) a speedy but incremental decline in the emergence of new specializations. For each component, I describe its realization first in biological contexts, then in scientific contexts. We will attend to the philosophical aspects of these phenomena later in the chapter.

Chapter 4 sketched the sort of change that occurs in both scientific and biological contexts when the set of factors to which a population is adapting remains fixed. The dominant sort of evolutionary change in these contexts is that of *refinement* of existing practices. We then looked at scenarios in chapter 5 where the set of selective factors undergoes a modest shift. With that shift comes a correspondingly modest change in the practices that undergo adaptation, one that can eventually lead to significant divergence in the practices of different groups.

I now want to look at cases in which the set of selective factors undergoes an abrupt, high-magnitude change. In biological contexts, these events

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

are witness to the simultaneous extinction of geographically widespread, ecologically disparate, and effectively unrelated genera—for example, the end-Permian event, in which roughly 96% of species were killed, or the end-Cretaceous event, in which roughly 75% of species were killed. The bolide that killed the dinosaurs at the end of the Cretaceous was about 6 miles wide, ejecting a cloud of dust and debris that darkened the sky for up to two years after impact (Bardeen et al. 2017). At the end of the Permian period, a region the size of Mexico was covered in a sheet of lava nearly 1 mile deep for about a million years, poisoning Earth's air and oceans (Erwin 2006). In general, extinction regimes that kill most things on Earth—regimes that generate a "crisis in the history of life"—require an environmental perturbation of an extremely high magnitude, and its influence needs to be able to reach across most of the planet's habitats. These perturbations are truly catastrophic.

A sufficiently high-magnitude, sufficiently widely distributed perturbation can directly result in death of enough victims to threaten the entire biosphere. By killing a sufficient number of taxa through its direct effects, a catastrophic perturbation can trigger the collapse of food webs around the globe, resulting in an even greater number of extinctions and fundamentally undermining each ecosystem's ability to remake itself in the image of its predecessor. This abrupt emptying of ecospace is what allows a mass extinction to play its unique role as the engine of high-magnitude evolutionary change—in particular, its tendencies to promote the emergence of evolutionary novelties and to fundamentally restructure ecologies on a global scale (Jablonski 2005).

To see how mass extinctions usher in evolutionary novelties, let us start by imagining a densely packed ecosystem, where the environmental conditions have persisted more or less unchanged for an evolutionarily appreciable amount of time. Because the same selective conditions have existed for so long, we expect most of the populations in this ecosystem to be nearly optimized for their niches, and we expect most niches to be occupied.

If a given population is near an optimum, major morphological changes are almost certain to negatively impact fitness (see chapter 4). Thus, under these conditions, we expect either minor changes in the direction of the optimum or minor changes orthogonal to the optimum, which allow some members to take advantage of an unexploited resource or unoccupied niche. These latter sorts of changes are the kind that facilitate ecological speciation, as discussed in chapter 5. In general, then, under evolutionarily long-standing

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

environmental conditions, natural selection is expected to be very strong and thus to permit only relatively small morphological changes.

Let us move on from here to a situation in which a single species goes extinct in a densely packed ecosystem that has been subject to long-term environmental stability. Like other members of the ecosystem, this unfortunate species (we may assume) occupied a specific ecological role, whose vacancy will be filled by whichever taxon or taxa respond most quickly and comprehensively. Here, the evolutionary dynamics do not differ fundamentally from those of the previous scenario in which a new species arose for ecological reasons. The loss of a single species will have only a modest effect on the ecological density, with the result that the new niche occupants are overwhelmingly likely to be minor variations on adjacent taxa that are themselves near a phenotypic optimum. Thus, for reasons similar to those described above, we do not expect the removal of a single species to foster the emergence of major morphological changes. That species' extinction will generally not clear enough ecological space to qualitatively change the selective dynamics in the region.

By contrast, mass extinctions often involve the clearing of entire *adaptive zones*, a term coined by G. G. Simpson to denote a "set of ecological opportunities that may be exploited by appropriately adapted species."[4] The expulsion of all the occupants of all the niches in an adaptive zone leaves a wide variety of ecological opportunities available to any taxa that might be able to occupy them. Under these conditions, occupying a new niche is a far less challenging affair than when an ecosystem is densely packed and already near its carrying capacity.

In a nearly saturated ecosystem, the competition to occupy an empty niche will be extreme. Victory will be given to whomever can infiltrate the niche and quickly monopolize its resources. The best candidates for this job will be those who are already very well suited for its conditions, enabling them to gain entry while still maintaining a high level of fitness. Their high level of fitness in the new niche allows them to quickly reproduce, generating enough members to monopolize the new resource in short order. Now, the candidates that are already very well suited for the new conditions are most likely to be those that occupy a very similar niche and can thus move into the new niche through a relatively minor, high-probability morphological change (see the transition from grape-eating to toast-eating discussed in chapter 5).

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

When an entire adaptive zone is emptied, however, the competition to occupy any of its niches bottoms out. In stable, saturated ecosystems, severe competition favors the evolution of increasingly refined solutions to a stable, well-defined set of challenges. But when the pool of inhabitants has been cleared by a mass extinction, no such competition exists, and a candidate need not be already very well suited for a new niche's conditions. Indeed, at this stage, a *niche*—in the normal sense of that word—exists only in the most general sense (e.g., terrestrial, marine, arboreal). Diversification proceeds rapidly, resulting in a lack of the kind of environmental stability that prevails in background times and within which Darwinian evolutionary change predominates. Under these conditions, the evolutionary process will not necessarily favor those candidates that can outcompete the locals by making a minor morphological change. A period of extreme non-Darwinian change has begun.

Because nature's tendency to discriminate against most forms of life has become extremely relaxed, the probability of making a large leap across morphospace is significantly raised. This result will tend to hold even as niches become increasingly well defined. Even though a candidate from a distant location in ecospace might have relatively low fitness compared with the new niche's phenotypic optimum, the absence of competition means that there is no need to rapidly saturate the new niche and corner its resources. Poorly suited candidates are not generally excluded from empty niches in empty adaptive zones the way they are in densely packed ecosystems. In such circumstances, large-magnitude variations have a much higher probability of success than they otherwise would. Here is how Erwin and Valentine describe the situation:

> The length of the jump is chosen from a probability distribution in which small jumps are common but increasingly longer jumps are increasingly rare, both because the farther the jump, the more different the new niche, and because it is deemed more difficult to adapt to larger changes than to smaller ones. This expectation leads to one more rule: the longer the jump, the larger the unoccupied area required for successful landing. In establishing a very different adaptive type, longer-jump lineages are required to colonize vacant adaptive regions because their adaptations would be far from perfect in their new niches and they need to avoid too much competition or other sources of interference. (Erwin and Valentine 2013, 229)

It is believed that the large morphological leaps made possible by empty adaptive zones are the foundation for the emergence of major evolutionary

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

novelties. One of the clearest illustrations of this process comes from comparing the early-Paleozoic diversification to the Mesozoic-era diversification. Both intervals lasted for roughly the same amount of time (182 MY and 183 MY, respectively). Both intervals experienced similar rates of diversification at the family level. But the early Paleozoic saw the emergence of many new phyla and classes, whereas the Mesozoic produced none. Erwin, Valentine, and Sepkoski (1987) attributed this discrepancy to the stark differences in ecological opportunities that were available during these times.[5] The vast amount of empty adaptive space available during the early Paleozoic enabled the kinds of high-magnitude morphological shifts that define higher taxonomic phenomena like the emergence of classes and phyla. By contrast, the Mesozoic diversification proceeds from a point at which most adaptive zones are already occupied, albeit at much lower densities due to the end-Permian extinction event. This and subsequent studies suggest that the vacancy of adaptive zones is the key ingredient in the emergence of major evolutionary novelty. Mass extinctions are the only known mechanisms capable of clearing adaptive zones.

The fossil record shows that during these times, an enormous range of highly disparate morphologies tends to appear. This trend of diversification and morphological divergence continues for the "(geologically) brief" period during which the normal "rules for extinction and survival" are altered (Jablonski 2005, 203). As nature experiments with a variety of different kinds of organization, some kinds leave more descendants than others—sometimes for Darwinian reasons, sometimes by chance. A positive feedback loop ensues, in which the increasing abundance of life provides more resources for yet further development (Erwin and Valentine 2013, 9).

As life on Earth begins to approach its carrying capacity, the pace of diversification starts to level off and change returns to a predominantly Darwinian mode. With ecospace now reaching its saturation point, environmental conditions become relatively stable, and lineages are once again rewarded mostly for increasingly small, increasingly refined solutions to increasingly well-defined environmental challenges. But although natural selection is again the paramount evolutionary power, there has been a radical reorganization of ecosystems made possible by the sudden disappearance of entire species, genera, families, and orders. When a catastrophic perturbation kills a taxon—especially a higher-level taxon—niche space is emptied, and the surviving taxa pour into the ecological vacuum. This

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

will often result in the rapid ascendency of previously marginal taxa, as was the case with mammals after the extinction of the dinosaurs. Prior to the Cretaceous-Paleogene extinction, mammals were a low-diversity, runty little order mostly confined to subsistence on bugs, worms, and other small, easily obtainable sources of nutrition that would have been too insignificant for a dinosaur to bother with. With the dinosaurs out of the way, mammals rapidly diversified, with concomitant increases in body size and population size.

In this way, mass extinctions result in an upheaval of the ecological status quo. Once the period of diversification begins to slow down, we do not see reversals of fortune for higher-level taxa; those that become dominant stay dominant. The catastrophic perturbations emblematic of mass extinctions seem to be the only thing that can dislodge incumbent taxa. Their disappearance provides an opportunity to fundamentally reorganize ecospace, the result of which is a new set of "rules" governing what it takes to be an evolutionarily successful clade (Jablonski 2005).

Although events like bolide impacts and lava are not governed by the evolutionary process, the biological response to such events *is*, even if that response is not thoroughly Darwinian in nature. According to the argument of this book, then, when the development of scientific knowledge is governed by the evolutionary process, scientific knowledge ought to exhibit the kind of signature phenomena when subject to the catastrophic perturbations that have appeared during periods of "crisis in the history of life"—namely, the emergence of fundamentally novel types of problems and solutions and a radical change in the kinds of approaches to problem-solving that tend to be successful.

6.2 Crisis and the Emergence of Scientific Theories

Kuhn was part of a generation of philosophers of science who were keenly interested in the kind of epistemically perturbative situation in which the state of research changes in such a way that, eventually, neither a refinement of an existing model nor an ad hoc addition to the family of models seems worth the trouble. This situation, which Kuhn describes as a moment of "crisis," emerges when our knowledge of the kinds of phenomena around which a particular specialization had initially developed has changed in such a way that we cannot bridge the gap between what we know on the

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

one hand and the sorts of inquiry that the specialization seems capable of directing on the other, quickly and satisfyingly enough to prevent the specialization's obsolescence. A response more radical than refinement or ad hoc adjustment is called for.

Now, the circumstances under which crises emerge vary enormously. But under this variety lies a key ingredient—namely, that the specialization no longer appears to be capable of efficiently solving the kinds of problems whose solution is in demand. Kuhn (1962b, 92) describes this state of affairs in a famous comparison between scientific and political change:

> Political revolutions are inaugurated by a growing sense, often restricted to a segment of the political community, that existing institutions have ceased adequately to meet the problems posed by an environment that they have in part created. In much the same way, scientific revolutions are inaugurated by a growing sense, again often restricted to a narrow subdivision of the scientific community, that an existing paradigm has ceased to function adequately in the exploration of an aspect of nature to which that paradigm itself had previously led the way. In both political and scientific development the sense of malfunction that can lead to crisis is prerequisite to revolution. . . . [T]hat parallelism holds not only for the major paradigm changes, like those attributable to Copernicus and Lavoisier, but also for the far smaller ones associated with the assimilation of a new sort of phenomenon, like oxygen or X-rays.

Once the specialization is perceived as being incapable of effectively occupying an available intellectual niche, it is marked for obsolescence—a "dead clade walking" (Jablonski 2002).

Kuhn's description of periods of "crisis" strongly resembles the states of affairs that prevail during mass extinction events. Particularly emblematic for Kuhn were two phenomena: the failure of existing practices and the proliferation of new and increasingly different forms of practice:

> The early attacks upon the resistant problem will have followed the paradigm rules quite closely. But with continuing resistance, more and more of the attacks upon it will have involved some minor or not so minor articulations of the paradigm, no two of them quite alike, each partially successful, but none sufficiently so to be accepted as a paradigm by the group. (Kuhn 1962b, 83)

The scientific community, which in mature instances guards the cluster of acceptable practices with admirable vigilance, is now forced to choose between preserving the character (if not the content) of consensus practice and giving into to their crisis-induced desperation by loosening the constraints on what forms of practice might count as "acceptable" after

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

all. "All crises," he writes, "begin with the blurring of a paradigm and the consequent loosening of the rules for normal research," suggesting that the commitment to a particular paradigm is only as strong as practitioners' ability to ignore a crisis (Kuhn 1962b, 84). "As in manufacture so in science—retooling is an extravagance to be reserved for the occasion that demands it. The significance of crises is the indication they provide that an occasion for retooling has arrived" (76).

With the rules of normal research relaxed, a period of "extraordinary research" begins and practitioners are permitted to explore a wide range of alternative problem-solving approaches in their efforts to bring nature to heel. In both its character and its trajectory this period closely resembles the recovery period following a mass extinction. As to character, research efforts during this time often involve a unique tolerance for philosophical reflection in the context of scientific inquiry proper:

> It is, I think, particularly in periods of acknowledged crisis that scientists have turned to philosophical analysis as a device for unlocking the riddles of their field. Scientists have not generally needed or wanted to be philosophers. Indeed, normal science usually holds creative philosophy at arm's length, and probably for good reasons. . . . [T]he search for assumptions . . . [can] be an effective way to weaken the grip of a tradition upon the mind and to suggest the basis for a new one. It is no accident that the emergence of Newtonian physics in the seventeenth century and of relativity and quantum mechanics in the twentieth should have been both preceded and companied by fundamental philosophical analyses of the contemporary research tradition. (Kuhn 1962b, 88)

Similarly, it is only during mass extinction events that we see the overturning of deeply entrenched aspects of morphology, such as body plans. As far as we know, these kinds of changes are impossible during background times, because the unrelenting pressure of natural selection will not allow for as wide and as deep an exploration of morphospace as is required to fundamentally change the basic morphological structures around which life is organized. Likewise, the premium placed on the refinement of existing scientific practices as a means of making progress tends to result in the exclusion of philosophical analysis as an unprofitable "extravagance." But desperate times call for desperate measures, and the failure of normal scientific procedures can result in a critical examination of existing practices that goes far deeper than anyone would find reason to undertake outside of a crisis.

Another "extraordinary" aspect of Kuhnian crisis periods is the "random sort of research produced" as practitioners grope wildly across practice space

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

in an attempt to gain a foothold in the new research environment (Kuhn 1962b, 87). This corresponds to the highly experimental nature of morphology that appears during recovery periods, as do its consequences: In both contexts, the failure of established means of making progress, combined with the consequent decline in the stability of the set of problems that demand solution and of problem-solving strategies that work best, results in a vast proliferation of new types of problems and of new types of problem-solving.

As in mass extinction events, the failure of incumbent practices in a given field of inquiry can often open up a variety of opportunities for practitioners to make progress. With the playing field now significantly leveled, some of the "divergent . . . numerous partial solutions that concerted attention to the [crisis-inducing] problem has made available" emerge as relatively more capable of structuring the ignorance into which practitioners have been plunged (Kuhn 1962b, 83). The eventual result is "a reconstruction of the field from new fundamentals, a reconstructing that changes some of the field's most elementary theoretical generalizations as well as many of its paradigm methods and applications" (85). In this closely analogous sense, Kuhn seems to have agreed precisely with mass extinction theory in holding that fundamental "novelty . . . can only emerge" with "the rejection of an older paradigm" (95–96).

In the previous section we discussed the other of mass extinction's two signature phenomena: the upheaval of the ecological status quo. We are now in a position to see the direct sense in which this phenomenon is sometimes realized in scientific contexts. Just as early successful biological lineages set the tone for what is evolutionarily possible in the post–mass extinction era, the novel scientific practices that are able to gain an early foothold during a crisis will ascend in due course to the rank of new normative constraints on which scientific practices will be permitted to contribute to the production of knowledge. It is interesting to note that Kuhn would later change his term for these constraints from the unwieldy "paradigm" to the more evocative "disciplinary matrix," a term that better captures the kind of reorganized disciplinary ecology that distinguishes the postrevolutionary period.

The sense in which disciplinary ecologies get reorganized is twofold. First, we often see a fundamental reversal in the kinds of practices that are emblematic of good science. My favorite example of this is the rapid transition from mechanistic theories to field theories in the seventeenth

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

and eighteenth centuries. The nonmechanistic commitment required by the Newtonian theory of gravitation was a serious defect from the perspective of many natural philosophers, and Newton was by no means insensitive to it (Janiak 2007). A generation later, Newton's theory dominated the intellectual landscape, and non-mechanistic field theories would continue to emerge for the next several generations. The regularity of this sort of reversal has been widely accepted since before Kuhn's work first appeared and should require no further elaboration.

Other ways in which biological ecologies get reorganized after a mass extinction involve the reoccupation of adaptive zones and the reconfiguration of niche space, both of which can be analyzed in terms of *ecological functions*. In your average, everyday ecosystem collapse, the species in charge of the various ecological functions—apex predators, keystone species, and the like—are wiped out. Interestingly, however, their ecological stations tend to be repopulated by members of the same higher taxon (say, a genus, or a family) drawn from the regional species pool. The only conditions under which this phenomenon tends not to occur are when mass extinctions clear out adaptive zones. Then—and only then, so far as we can tell—can there be radical turnover in the higher-order taxa occupying these functional roles in the ecosystem. At this time, we also see the radical reconfiguration of niche space. New and fundamentally different kinds of niches emerge, providing hitherto unavailable opportunities for adaptive divergence. Relatedly, adaptive zones will be repopulated by different higher-order taxa, such as the repopulation of the skies by birds after the pterosaurs disappeared following the Cretaceous-Paleogene extinction.

The radical reversal of functional roles, the emergence of entirely novel functional roles, and the functional recategorization of existing practices are at the core of what Kuhn took to be the defining characteristics of a major scientific revolution. The conviction with which he held this view, and the clarity with which he expressed it, would only increase with time:

> The distinctive character of revolutionary change . . . is that it alters not only the criteria by which terms attach to nature but also, massively the set of objects or situations to which those terms attach. What had been paradigmatic examples of motion for Aristotle—acorn to oak or sickness to health—were not motions at all for Newton. *In the transition, a natural family ceased to be natural*; its members were redistributed among preexisting sets, and only one of them continued to bear the old name. (Kuhn 2000, 29–30; emphasis added)

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

And, like biological contexts, the interdependence of practices means that changes that occur at very fundamental levels—for example, changes in the meaning of the term "motion"—cannot but disrupt the delicate web of relations that characterizes the system: "since such redistribution always involves more than one category and since those categories are interdefined, this sort of alternation is necessarily holistic" (Kuhn 2000, 30).

Our goal in these past couple of sections has been to see how closely mass extinctions resemble Kuhn's image of major scientific revolutions, both in the historical pattern of the phenomena that characterize them, and in the causes that we now think account for those phenomena. With respect to both kinds of events, I have tried to balance the descriptive brevity and generality that is warranted by their somewhat peripheral role in the book's overall argument on the one hand with the grain of specificity that is required to convincingly argue that the incredibly close analogy between the two is no mere artifact of vacuously general description on the other. I hope that I have succeeded.

For any holdouts, there are additional means of persuasion we can employ that appeal not to analogous empirical patterns per se but to various philosophical implications that the analogy holds for the epistemology of scientific knowledge. In ways of which I never conceived at the outset of this project, the contrast between normal science / adaptive optimization on the one hand and scientific revolutions/mass extinctions on the other allow for a full appreciation of the epistemological significance of evolutionary theory. For it is only with this contrast firmly established that it becomes possible to understand the fundamentally evolutionary character of core epistemological concepts of *progress*, *success*, and *rationality*. As an added bonus, we will see clearly—more clearly, I think, than has previously been possible—why Kuhn never abandoned the view that science makes progress of an epistemically respectable variety during both normal and revolutionary times.

6.3 The Epistemic Significance of Non-Darwinian Scientific Change

6.3.1 Non-Darwinian Populations

Recall this book's overarching thesis—that an important part of the epistemic power of science derives from the role that natural selection plays in

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

the modification of scientific practice. To support this thesis, I have argued that the degree to which scientific practice satisfies the criteria for being a paradigmatically Darwinian population tracks the degree to which it is able to generate the two forms of epistemic progress that we consider to signatures of the epistemic power of science—namely, (1) its tendency to develop ever more refined versions of problems and their solutions (vertical progress) and (2) its tendency to grow (horizontal progress).

I'd like to begin our philosophical investigation of the parallels between revolutionary scientific change and revolutionary biological change by first recasting the previous sections' discussion of non-Darwinian change in the language of Darwinian populations. Ideally, we should be able to explain instances of non-Darwinian change in terms of a population's failure to satisfy certain of the criteria for being a Darwinian population.[6] Therefore, let us now look at which of the relevant criteria are compromised during a moment of "crisis."

I first attempted to articulate the sense in which crisis-period change is non-Darwinian by referring to the fact that relative *adaptedness* seems to play an increasingly marginal role during such times, as other, extrinsic factors begin to gain influence. Here, the kinds of properties that give populations an edge are those that tend to provide an advantage when the sampling process becomes random—factors like high relative reproductive rate, high relative intragroup variation, and wide geographic range. High reproductive rate increases the probability of being sampled at a given location, whereas wide geographic range increases the probability of being sampled regardless of location. High intragroup variation increases the probability that a member of the sampled population will be better adapted to changing environmental conditions. Because of the predominant importance of these extrinsic factors, populations plagued by crisis tend to receive low marks on the *intrinsicality* (*s*) dimension of Darwinian populationhood. Thus, one symptom of crisis is a generalized failure of intrinsic properties to explain evolutionary change.

Although Kuhn did not, as far as I know, use the language of "intrinsicality," his discussions of crisis periods consistently refer to "loosening of the rules of normal research" (which in the present study have become equated with the rules of adaptedness) (Kuhn 1962b, 84). During a crisis, success tends to be less about how well this or that theory explains this or that phenomenon, as Kuhn famously observed: "paradigm debates are not really about relative problem-solving ability" (157). "Instead," he continues,

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> the issue is which paradigm should in the future guide research on problems many of which neither competitor can yet claim to resolve completely. A decision between alternate ways of practicing science is called for, and in the circumstances that decision must be based less on past achievement than on future promise. (157–158)

In contrast to scientific judgments made in the course of normal science, choosing between alternate ways of practicing science itself are made "for highly individual reasons," based on "arguments that appeal to the individual's sense of the appropriate or the aesthetic." In these debates, "the new theory is said to be 'neater,' 'more suitable,' or 'simpler.'"[7] In general, decisions made in these circumstances tend to be more subjective—that is, they have more to do with the preferences of individual scientists—than under conditions of normal science, where objective relations between theory and evidence predominate. Recast in the language of "intrinsicality," scientific judgments made in crisis mode have to do less with the intrinsic features of scientific practice and more to do with extrinsic factors, such as an individual's particular aesthetic sensibilities.

The appearance of philosophical reflection on practice is another example of the increasing influence of extrinsic factors during crisis. Philosophy is appealed to precisely at times when the means of scientific judgment that have been developed and validated internally begin to fail. Philosophy—always regarded by scientists as outside science (sometimes to its credit, sometimes not)—is brought in to provide some structure to an area of research whose once well-ordered relations have collapsed; it will be evicted just as abruptly once order has been restored.

Kuhn repeatedly emphasized the importance of potential *fruitfulness* as a criterion for assessing a new paradigm's "future promise." Fruitfulness in science is a difficult notion to cash out satisfactorily, but it is closely aligned with the ensemble-level properties mentioned above—high relative reproductive rate, high relative intragroup variation, and wide geographic range. A fruitful research framework will give more insights (reproductive rate) into more kinds of problems (wide geographic range), which it does by being well suited for lots of different challenges (intragroup variation).[8] Just as in mass extinction events, these ensemble-level properties had, for Kuhn, particularly strong influence over the development of scientific knowledge during periods of crisis.

Each of these types of factors—individualistic/aesthetic, philosophical, and ensemble-level—contributes independently to the low *intrinsicality*

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

score that populations of practices receive during crisis periods. This is not to say that problem-solving ability/adaptedness is inoperative or irrelevant during a crisis—it is never inoperative. But its relative importance is significantly diminished, and with that comes the associated rise in the importance of factors that bear far less directly on problem solving. Those factors, too, are present in times of normal science, but their influence is overshadowed by the paramount power of selection to modify populations when *intrinsicality* is high.[9]

Another Darwinian population criterion we described in chapter 1 was *continuity*, a property of fitness landscapes that facilitates adaptive modification by situating morphotypes with similar fitness values adjacent to one another in morphospace. In a landscape with high continuity, small changes in morphology (small steps in morphospace) result in slight changes in fitness. Because small changes in morphology are more probable and the probability of deleteriousness increases with the size of the change, selection is able to push populations on highly continuous landscapes "smoothly" toward an optimum. By contrast, optimization on discontinuous landscapes is hampered by a situation in which adjacent locales in morphospace tend to differ significantly in how fit they are. Under these conditions, it is no longer the case that the probability of deleteriousness necessarily increases with the size of the change in morphology; a particular point in morphospace might be surrounded by deleterious steps, for example. Godfrey-Smith similarly describes a kind of situation such that "when we have an extremely low value of *C*, we have a situation with an almost chaotic character. Each total set of properties of an individual gives rise to a particular reproductive output, but a very slight modification of those features will have a different upshot" (Godfrey-Smith 2009, 62).

This is an apt description of fitness landscapes during mass extinction events. In contrast to the smooth adaptive ascent that is characteristic of landscapes in background times, the radical instability of the environment in periods of crisis produces fitness landscapes that are not very well defined topologically—and, a fortiori, not very continuous. A fitness landscape specifies the fitness of a morphotype in a given environment. When environmental conditions undergo perpetual, high-magnitude fluctuations, they lack the stability required for a population to undergo adaptive optimization. To put this in terms of high or low *continuity*, unstable environmental conditions have extremely low values of *C*, because the magnitude of a morphological change is not a good predictor of the magnitude

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

of a fitness change. The most evolutionarily significant consequence of the lack of environmental stability is that whatever is fit today might be unfit tomorrow. And because changes are likely to be small, a population cannot hope with high probability for a positive change in fitness. When selective factors are in flux, there cannot be adaptation.

In a sense that should now be easy to see, destabilized conditions are a classic element of Kuhnian crises. Indeed, Kuhn frequently describes paradigm debates precisely in terms of debates over which problems have been solved or need to be solved, which methods are effective, and which species of evidence are probative. During a period of crisis, the epistemic landscape lacks the high degree of continuity required for the rapid rate problem-solving that occurs during normal times. With the "loosening of the rules of normal science," the stability of the set of practices can no longer be counted on, and the kind of refinement of problems and methods that characterizes normal science is impossible. In this kind of scenario, scientists behave (according to Kuhn) exactly as we would expect them to, with progress in the near future seeming to depend less on understanding this or that bit of nature and much more on imposing some structure or other on our ability to ask and answer questions.

The last problem posed by scientific crisis for Darwinian populations lies along the *heredity* (*h*) dimension. In chapter 2 we looked at the many ways that scientific training aims at producing high-fidelity copies of scientific practice from one generation to the next and noted how well this would suit any endeavor whose main goal was solving problems. Kuhn, keenly aware of the epistemic function of scientific education, contrasts this normal-science tendency with crisis situations, in which initiates to a given field show far less fidelity to consensus practice than did generations reared during simpler times:

> Almost always the men who achieve these fundamental inventions of a new paradigm have been either very young or very new to the field whose paradigm they change. And perhaps that point need not have been made explicit, for obviously these are the men who, being little committed by prior practice to the traditional rules of normal science, are particularly likely to see that those rules no longer define a playable game and to conceive another set that can replace them. (Kuhn 1962b, 90)

To summarize the argument to this point: one test of the concept of a Darwinian population is whether it can capture the characteristic patterns

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

of change during a mass-extinction event, during which the non-Darwinian aspects of evolutionary change tend to be abnormally influential. We can accurately describe the characteristically non-Darwinian sort of change that occurs during mass-extinction events in terms of low values of *s*, *h*, and *C*—that is, those quantities that seem to matter most to a population's susceptibility to adaptive modification. We can with similar success describe the sort of scientific change that occurs during Kuhnian crises, which I have alleged to be the analogue of the "crises in the history of life" in which so much non-Darwinian change occurs. Accordingly, we can say that the reason that scientific change is so abnormal during periods of crisis might be because scientific practice comes to possess low values of *s*, *h*, and *C*, thus compromising the degree to which it approximates a paradigmatically Darwinian population.

6.3.2 Progress through Revolutions

The upshot of this discussion is that my argument seems to commit me to the view that, in short, a lot goes on during a major scientific revolution that does not contribute to the special epistemic status of science. If (1) a science's ability to make epistemic progress varies directly with the extent to which it approximates a paradigmatically Darwinian population, and (2) science scores low on the Darwinian population scale in several important respects during a scientific revolution, then (3) it must follow that scientific revolutions are, at least in some respects, epistemically otiose.

Let me be even more frank. Earlier I argued that the most distinguishing features of a scientific revolution—fundamentally novel ideas and new norms governing inquiry—derive specifically from the non-Darwinian character of revolutionary scientific change. Doesn't that mean that fundamentally novel ideas and new norms governing inquiry can only emerge when scientific practice fails to satisfy in a significant way the constraints on being a Darwinian population—that is, only when scientific practice is least apt to make an epistemically progressive contribution to human knowledge? Aren't radical new ideas and new norms of inquiry part of the distinctive epistemic power of science?

I want to begin to answer these questions by going back to something I quoted earlier that Kuhn says in the final pages of *Structure*: "novelty for its own sake is not a desideratum in the sciences as it is in so many other

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

creative fields" (1962b, 169). Earlier I used this sentiment as part of the Darwinian explanation for why some disciplines are very good at solving problems—namely, that resistance to novelty contributes to high values of *heredity* and *intrinsicality*. I now want to elaborate on this sentiment as part of an explanation for why the production of radical new ideas per se is not part of what gives science its distinctive brand of epistemic power. I will then turn to the question of whether transitions in norms of inquiry per se contributes to the epistemic power of science.

There is no question that novel ideas are of immense value in scientific inquiry. Virtually all of the epistemically salutary phenomena I've described in this book require novel ideas, and it is essential to any recognizable notion of scientific progress. But that is not to say that novelty per se is valued in science. Rather, it is *necessary* novelty that is consistently in high demand. Novel ideas that do not address some pressing need in the sciences are often actively resisted, for what I think are unassailably good epistemic reasons. Kuhn understood this point well:

> It is often said that if Greek science had been less deductive and less ridden by dogma, heliocentric astronomy might have begun its development eighteen centuries earlier than it did. But that is to ignore all historical context. When Aristarchus' suggestion was made, the vastly more reasonable geocentric system had no needs that a heliocentric system might even conceivably have fulfilled. (Kuhn 1962b, 75)

A short while later he tries to spell out what gives rise to the kind of need that novelty alone can fulfill:

> The scientist who pauses to examine every anomaly he notes will seldom get significant work done. We therefore have to ask what it is that makes an anomaly seem worth concerted scrutiny. . . . Sometimes an anomaly will clearly call into question explicit and fundamental generalizations of the paradigm. . . . Or . . . an anomaly without apparent fundamental import may evoke crisis if the applications that it inhibits have a particular practical importance. . . . Or . . . the development of normal science may transform an anomaly that had previously been only a vexation into a source of crisis. (Kuhn 1962b, 82)

Whatever the proximate cause of crisis, for Kuhn, the most "important source of change is the divergent nature of the numerous partial solutions that concerted attention to the problem has made available" (1962b, 83). These varied attempts to fulfill the need created by an anomaly are often the predecessors of the proliferation of specialties that Kuhn took to be an essential component of scientific revolutions and one of the two senses in

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

which science makes epistemic progress.[10] I presented my own long evolutionary argument for this view in chapter 5, in which I claimed that the emergence of new specializations is epistemically driven when it is fueled by the availability of a significant unsolved problem.

The relevant question at this point is whether the novel specializations that emerge in the wake of scientific revolutions are fundamentally different in kind from what emerges during a run-of-the-mill bout of ecological divergence. I think that the answer to this question is a well-established "no." The difference is one of degree, not of kind. First, not all scientific revolutions exhibit radically divergent developments; this only appears to be the case for large-scale revolutions. The discovery of X-rays, for example—an exemplar of a revolution for Kuhn[11]—involved no radical divergence. The transition to general relativity did. So the radical degree of divergence is not a necessary consequence of scientific revolutions per se. In general, the degree of divergence should be proportional to the size of the crisis—that is, to the volume of niche space that has been cleared. For some revolutions, the magnitude of that volume will be profound, and the resulting degree of divergence will be correspondingly huge, as it is with mass extinctions that clear entire adaptive zones (or, for that matter, in the case of the Cambrian explosion, which could be described as a transition from pre-paradigm life to a paradigm-governed life).

If we accept that the degree of divergence is proportional to the size of the crisis, then radical divergence per se is not an inevitable consequence of revolutionary transitions. But divergence itself is an inevitable consequence, because revolutionary transitions are at their core concerned with the pursuit of qualitatively novel problems. In large-scale transitions, it just happens to be that there are many such problems. The competition to occupy the space created by any particular unsolved problem will thus be correspondingly low, permitting the sort of high-magnitude divergence we expect in biological contexts. The radical nature of the divergence is the non-Darwinian element of the proliferation of specializations that occurs during large-scale revolutions. But the divergence itself is Darwinian, and it is driven by the same Darwinian mechanism in both large-scale revolutions and low-level branching events.

These results comport nicely with the Kuhnian view that novelty is not valued for its own sake. More importantly, it preserves Kuhn's view—which would only gain in strength after *Structure*—that the epistemically progressive contribution of scientific revolutions is the increase in specializations.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

It is this contribution that is an inevitable consequence of the fact that scientific practice partly maintains its status as a Darwinian population during times of crisis; perhaps it is better conceived of as a Darwinian *metapopulation* under such conditions.

The non-Darwinian origin of the new norms for inquiry that emerge out of a crisis is less complicated epistemically. The shift in epistemic and metaphysical norms that Kuhn makes a big deal out of in *Structure* was unquestionably the most epistemically problematic feature of his model, as it seemed to undercut science's capacity for producing knowledge while giving aid and comfort to relativists. I am sympathetic to this view. Certainly there is nothing inherently progressive about changing the norms of inquiry; it might be progressive, or regressive, or maybe simply a lateral move. And this is precisely the result we should expect if the epistemically progressive sort of scientific change is Darwinian change. We can have no idea whether non-Darwinian scientific change will be epistemically progressive or not. Insofar as the changing norms of inquiry—"paradigm shifts"—are a consequence of non-Darwinian change, there is no reason to expect that the practice of inquiry itself will be epistemically better off. Of course, this not to say that changing norms of inquiry cannot be progressive. Indeed, the norms of inquiry are continuously refined through the same Darwinian process that produces the adaptive refinement of lower-level components of practice like theories and problems. The steady evolution of means of data collection and analysis since the seventeenth century, about which much has been written, is a clear example of this process.

To summarize: The non-Darwinian aspects of some revolutionary changes in the history of science can be understood in terms of the failure of scientific practice in those cases to satisfy to a sufficient degree the criteria for being a Darwinian population. But these aspects—radical divergence, and new norms of inquiry—do not by their very nature contribute to the epistemic power of science. *Divergence* does, but divergence itself is a Darwinian phenomenon, whether it takes place during normal science or crisis.

6.4 Revolutions, Incommensurability, and Non-Darwinian Change

6.4.1 Kuhn's Generalized Notion of Revolutions

Isolating the distinctive effects of scientific revolutions in this way allows us to shed light on Kuhn's infamous preoccupation with "the increasing difficulty of communication between the practitioners of different specialities"

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

or *incommensurability* (Kuhn 1993, 337). In particular, we can see how the view that scientific revolutions are distinctive by virtue of their causal role in the production of novelty and new norms of inquiry helps to explain Kuhn's sense of what counts as a revolution, and why all revolutions have the same attendant epistemic problems.

For Kuhn, each revolution involved a branching event in the history of science—the emergence of a new scientific specialization. Over time, he would increasingly refer to "the speciation-like process through which new disciplines emerge" (Kuhn 2000, 100). As we have seen, some evolutionary branching events are minimal from a taxonomic perspective, as in the case of ecological speciation described in chapter 5. Likewise, the emergence of a new scientific discipline can be an extremely local event, often involving no more than a handful of practitioners; this is particularly true in our era of hyperspecialization.

However, other evolutionary branching events are not so minor. They involve major taxonomic changes, such as the emergence of new classes, orders, and phyla. These events are the sort that only mass extinctions seem to be capable of generating. Similarly, we tend to see major disruptions in the organization of scientific inquiry only in during truly epic scientific revolutions. For example, it is widely held that modern physics did not exist as a scientific discipline before the Scientific Revolution; or, to use Kuhn's example discussed earlier, optics before Newton. In other words, "revolutions do differ in size and difficulty" (Kuhn 1993, 337). Kuhn was consistent about this from *Structure* until 1993's "Afterwords."

With that distinction in place, let us now tackle the more nebulous question of why "the epistemic problems [revolutions] present are . . . identical" (Kuhn 1993, 337). First, what are the epistemic problems? These are the famous Kuhnian problems—incommensurability, world changes, shifts of vision, and (later) lexical and taxonomic divergence. I want to treat these as different attempts to capture the same phenomenon or set of very closely related phenomena, but I believe Kuhn came to see the phenomenon of lexical divergence as basic. In keeping with Kuhn's usage I will structure my discussion around the notion of lexical divergence and I will flag instances where there is a need to specifically refer to one or another of the other attendant ideas.

Now, what reason is there to think that the epistemic problems are identical regardless of revolutionary size and difficulty? First, conceived of as branching events, it is almost beyond question that all revolutions will

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

involve some amount of divergence, for the same reason that speciation events will involve some amount of divergence. The isolation of a reproductive group—organisms or practices—will ensure at least some divergence between that group and its parent group. The degree of divergence will be directly proportional to time since divergence, the dimensionality of the state space, the degree of isolation, and the degree of difference in selective factors affecting the two groups. In addition, divergence will be exacerbated by small group size. So the fact that there will be *some* divergence in practices between the two groups seems unassailable, if we characterize revolutions as branching events.

Second, insofar as branching events involve some variety of isolation, that divergence in practice will be at least partly lexical. From the point at which branching occurs, terms used in the course of a group's scientific investigations will begin to undergo refinement in response to the specific suite of selective pressures that is unique to that group. As part of scientific practice, those terms are subject to the same kind of adaptive refinement that affects other dimensions of practice.[12] To the extent that a newly emergent group investigates novel problems, investigation of those problems will pull their lexicon in a particular direction, along with the rest of their practices. To the extent that the new group is isolated from other groups, their lexicon will diverge through drift, as will the rest of their practices. Thus, so long as we take revolutions to be branching events, I find the conclusion that every revolution involves lexical divergence to be basically unavoidable. In this sense, Kuhn was correct in thinking that all revolutions face the same epistemic problems.

Nevertheless, the size and dimensions of divergence will differ from case to case, in association with the scale of the branching event. But this too is a straightforward consequence of evolutionary theory. Starting from one end of the branching continuum, to whatever extent an ecologically generated species and its parent remain ecologically similar, their degree of morphological divergence will probably not be exceptional. Similarly, a minimally disruptive branching event in science will probably not immediately be associated with a radically different lexicon. In contrast, it seems only to be in the case of mass extinctions that we get radical morphological divergence in the history of life. In the same vein, the major revolutions in the history of science tend to involve correspondingly major changes to scientific practice. In addition, the fact that divergence can occur along

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

different *dimensions* of scientific practice suggests that we should not expect all instances of divergence to look the same, or perhaps even similar. Some instances may initially involve only a subtle, nearly imperceptible shift in visualization of a formalism, for example. Other, larger-scale instances may involve a marked shift in the standards for considering a puzzle to be solved, as in the revolutionary transitions that occurred in nineteenth-century mathematical proof theory.

The scale and direction dependency of lexical divergence, in turn, implies that many instances of lexical divergence may be either extremely subtle or otherwise difficult to detect. For example, to the extent that an ecologically generated specialization and its parent remain similar in the kinds of problems and methods they investigate, much of their lexical divergence will probably be too subtle to notice, perhaps even in very special circumstances. In order to detect lexical divergence at this subtle of a level, we may have to look for its downstream effects—as with the subsequent proliferation of discoveries that followed the identification of Uranus (Kuhn 1962b, 115–117)—rather than direct evidence that two communities don't understand each other. Plainly put, we simply should not expect every instance of lexical divergence to provoke mass confusion. (And let's recall that things are no different in biology. Even with gene sequencing, there is often still debate over whether two varieties constitute different species.) But during major revolutionary transitions, instances of lexical divergence will be relatively easier to spot: Newtonian "mass" and relativistic "mass" are obviously not equivalent notions. These are the cases that make Kuhn's claims about epistemic problems look most convincing.

6.4.2 Kuhn's Meaning-Incommensurability

I have shown that each revolution is expected to involve lexical divergence due to the evolutionary implications of construing revolutions as branching events. It is a short distance from this result to the conclusion that every revolution involves some amount of Kuhn's "meaning-incommensurability" (hereafter just "incommensurability").[13] Given how universally reviled Kuhn's incommensurability thesis is among philosophers of science, the ease with which my account accommodates incommensurability may rightly be perceived as a symptom of the evolutionary model's invalidity. Nevertheless, I think it is important to see whether we can preserve Kuhn's idea in as natural a way as possible, if for no other

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

reason than to validate his basic intuition that he himself had presented an evolutionary model of science.

Before deriving that conclusion, let's consider the following objection: even if it is true that all revolutions involve lexical divergence, it does not immediately follow that all revolutions generate incommensurability. For, it is prima facie possible that practice terms could be translated from the lexicon of one specialization into the lexicon of the other. And, indeed, we often (though by no mean always, or perhaps even generally) find that members of different specializations have no problems communicating with one another.

On this last point, we should not take even seamless communication as evidence for the absence of incommensurability. Effective communication requires something generally less demanding than lexical translation. For one thing, communication can benefit from contextual cues of which translation cannot avail itself. Communication can occur via many different channels, including diagrams that admit of no precise description but are nevertheless central to scientific practice. Any specialization that uses diagrams is thereby saddled with an inability to lexically translate, even though attempts to use diagrams to communicate across practitioners or across specializations may nonetheless still be successful.

Thus, when Kuhn says that "communication across the revolutionary divide is inevitably partial," it simply does not follow that communication will be problematic (Kuhn 1962b, 149). Partial communication is quite often good enough. A historically significant example of this platitude involves the interpretation of the adaptive landscape, a heuristic device used for diagrammatically representing fitness values for a given character state. This diagram was introduced by Sewall Wright in the 1930s, who used it to illustrate aspects of the mathematical theory of population genetics. These diagrams would have a profound effect on the great population geneticist Theodosius Dobzhansky, despite the fact that he was "literally innumerate," according to his student Richard Lewontin.[14] I would say that communication in this case was partial at best. Part of that incompleteness is due to Dobzhansky's innumeracy. But probably a larger part of it is due to the nonexistence of a mathematically precise description of the adaptive landscape itself. Nevertheless, Dobzhansky (and Ernst Mayr, and G. G. Simpson) got the essential spirit of Wright's diagrams and used those lessons to great effect (Provine 1971; Plutynski 2008).

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Generally speaking, we need to resist the temptation to descend into fits of epistemological angst when confronted with the notion that communication across specializations or across practitioners is "inevitably partial." The Dobzhansky example is one of scores that could be adduced to illustrate the productive epistemic power of diagrams that do not admit of numerical description, and thus can't be "translated" in any literal sense. Indeed, the lack of numerical description is arguably often an epistemic boon because it provides an opportunity for practitioners to think about the phenomena in a wide variety of ways that probably would not be available to them were they forced to think about things in purely algebraic terms.[15] So, although incommensurability might imply in*complete* communication, it does not imply ineffective communication, nor even mildly problematic communication. Barriers to communication will depend on the scale and the dimension of lexical divergence, not on the mere fact of untranslatability. Sometimes the scale of divergence will be so enormous that practitioners can't even communicate across a revolutionary divide without "going native." These cases are certainly fascinating, but they are not part of the essence of a branching event, even if incommensurability is.

If we avoid the misstep of invoking communication as evidence against incommensurability, the objection survives: what is inherent in revolutions that would stop practitioners from translating in principle? Kuhn's preferred entrée into incommensurability was a version of functional role semantics (Block 1986; see also Kuhn 1982). Scientific communities, he argued, are very specialized language communities. The terms of these languages get their meanings from those terms' roles in scientific practice. Different communities are defined by different sets of practices. Ergo, the same term (e.g., "mass") could not occupy precisely the same conceptual role in two different communities. Ergo no term can be translated from the language of one community into another.

This argument has always held a certain seductive appeal for me. It becomes more appealing when viewed from an evolutionary perspective, as Kuhn would come to appreciate (see quoted portion of Kuhn 1993 in section 3.2.2). There are many views on how species are individuated, of which Kuhn's is one, popular particularly within ecology. The standard mode of representation of ecosystems in contemporary ecology is that of a web, with species at the nodes and links between nodes representing

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

various sorts of ecological dependence. These links are the relations that make a difference to an organism's actual reproductive success. Significant changes to the environment equate to a significant reordering of relations between species, and between species and their abiotic environments. To the extent that a species can be partly characterized by its functional role in an ecosystem at a given time T_1, major environmental changes at T_2 translate to major revisions in the functional roles of that species. For Kuhn, theoretical terms are subject to the same sort of dynamic in semantic webs.

As attractive as I find this argument, it does seem to be sort of a cheap way to derive incommensurability, and the alleged correspondence between science and biology requires a bit more metaphor than I'm comfortable indulging in. For one thing, I bristle at the notion that "a species and its niche are interdefined." It is not scandalously problematic to say that a niche is defined in relation to a species. Many biologists hold the view that there are no niches out there in the world that exist independently of occupants. But it is mildly scandalous to say that a species is defined relative to its environment. Relations of descent are usually taken to be primary from a definitional perspective, although there is debate on this matter. So maybe it's just a personal failing that I read Kuhn's perspective as an attempt to use metaphor to skate over some nontrivial differences in the conceptual foundations of evolutionary biology. What is unexceptionable is that the notion that "a species and its niche are interdefined" is by no means unexceptionable.

The primary reasons that I won't appeal to this argument is that there is no explanatory role for the evolutionary process per se. Kuhn's meaning-incommensurability follows trivially from his functional role semantics. I think it would be preferable to derive the result as an effect of the evolutionary process, partly because it would be yet another illustration of the evolutionary model's power, but also because it is explanatory in a way that the appeal to functional role semantics is not. Had the development of scientific practice not been governed by the evolutionary process, incommensurability might be avoidable, even if meanings are fixed by terms' functional roles.

Each instance of branching, be it in scientific or biological contexts, involves the differential modification of "homologous structures" in response to different selective pressures, and in response to the effects of variation and population size (Kuhn 1982, 683). When branching occurs, the evolutionary

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

fate of the practices—including the "homologous structures"—of the newly emergent group is severed from those of its parent. As elements of scientific practice, the separation of evolutionary fate extends even to the meanings of terms. Kuhn, of course, would have allowed that meanings do not completely overlap even between two labs working contemporaneously. To the extent that such labs can be differentiated by slight differences in the set of scientific challenges to which each is responding, the evolutionary model suggests meanings-divergence even at this very local level. Large-scale scientific change, on the order of a major scientific revolution, tends to exhibit much more radical "world-view changing" behavior because the problem-solving tools are diverging in response to much larger-scale differences in the sets of problems they're being used to pursue.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

II The Emergence of a Discipline

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Up until now, this book has largely been an attempt to use evolutionary theory to secure a philosophical thesis about science: *Part of the immense epistemic power of science derives from the way in which the structure of scientific communities ensures that scientific knowledge will increase in depth and breadth*. The foundation of this guarantee is the susceptibility of scientific practice to the influence of natural selection. We know that scientific practice is susceptible to the influence of natural selection in many cases because scientific practice often approximates the conditions for being a paradigmatically Darwinian population.

As part of the effort to secure this philosophical thesis in a responsible and convincing manner, I have frequently had occasion to elaborate on the correspondences between the modification of scientific practice on the one hand and evolutionary biological change on the other. Scientific knowledge increases in depth, I have argued, because of natural selection's power for adaptive refinement of a population's fit with its environment. The process of adaptive refinement hones scientific practice to deliver increasingly precise answers to increasingly well-defined questions. I also have argued that the increasing breadth of scientific knowledge derives fundamentally from the power of natural selection to cause the branching of Darwinian populations: the process of adaptive divergence often drives scientific practice into the investigation of unexplored species of phenomena.

In both evolutionary and scientific contexts, each of these tendencies has been associated with a certain variety of progress. Fundamental to each variety of progress is the notion of *specialization*: increasing depth corresponds to increasing degrees of specialization, while increasing breadth amounts to increasing the number of specializations. Both of these tendencies are

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

immediately apparent when one examines either life or knowledge from a historical perspective; this is no accident. Both domains frequently show a drive toward increasing the degrees and kinds of specialization because both are strongly influenced by the causal power of natural selection, and that is just what natural selection does.

I now want to look in detail at the process of how new scientific Darwinian populations emerge. Chapter 5 was devoted largely to showing how a new research problem's ability to give rise to a new specialization depends on a couple of factors that feature importantly in the Darwinian process. First, it depends on whether the new research problem is distinct enough from existing research to constitute a new kind of research problem. Like an untapped environmental resource, the new research problem cannot be accessible to all members of the population if it is going to fuel the ecological divergence process. It must be the case that only certain members of the scientific community are equipped to profitably mine the depths of the new problem. At the same time, however, if research on the problem is going to endure long enough to make a lasting contribution to scientific knowledge, the significance of the problem must be recognizable even to those practitioners who are unable to investigate it themselves. Modern scientific inquiry requires lots of resources, but the pie is limited. In order to get a piece of that pie, a practitioner must be able to convince a very broad range of specialists that the research is worthwhile. The ideal situation for driving the emergence of a new specialization would be one in which there exists a subgroup of practitioners who were uniquely capable of investigation a problem that was widely recognized within the broader community as being of major importance. Unfortunately, most problems are not like this, because the two determining factors are constitutionally opposed to one another. Ceteris paribus, the more distant a new problem is from current research, the more difficult it is for practitioners to appreciate its significance.

Chapter 5 also stated a general principle for a discipline's propensity to "speciate": the probability that a new specialization will emerge varies inversely with the *marginal utility of adaptive progress in the parent discipline*. When it is easy to make major contributions to knowledge within an established discipline, the chances of that discipline "speciating" are low; those chances become better and better as it becomes more and more difficult to

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

make even a small contribution. Why? Because very small contributions to knowledge are plentiful. There is a veritable cornucopia of them, to most of which no established discipline can yet lay claim. And it is generally much easier to convince uninitiated colleagues of the minor significance of a problem than it is to convince them of its monumental importance. Thus, we should most expect the emergence of a new specialization under the following three conditions:

1. Most of the major contributions to a given specialization have been made.
2. A new problem has been identified whose significance is widely appreciated.
3. Investigating the new problem requires a specific set of investigative tools.

The purpose of part II is to carefully illustrate how these three conditions gave rise to the modern discipline of evolutionary paleontology, an event that possess many of the right features for testing the model of epistemic growth developed in this book. First, evolutionary paleontology's emergence as an autonomous discipline is recent and well documented, and so we have very good data on what happened and why. Second, evolutionary paleontology is relatively narrow in terms of the sources of evidence to which it appeals and the range of phenomena it attempts to explain, making it relatively easy to identify the key scientific issues involved. Third, the principal actors involved in evolutionary paleontology's gestation are few in number, further facilitating accurate identification of key scientific issues and making their relevance to influential individuals easier to track. Again, as I stressed in chapter 5, I do not claim that all new specializations must emerge in this way, only that it is the default expectation for paradigmatically Darwinian populations.

Our study follows the history of evolutionary paleontology through the lens of the storied career of Stephen Jay Gould. There is a simple reason for this: Gould was trying to establish paleontology as an autonomous evolutionary discipline during a time when there was widespread faith that all of the major contributions to evolutionary theory had already been made and that our understanding of the evolutionary process was more or less complete. He failed several times, and he failed for reasons that are straightforwardly predicted by the evolutionary model of scientific growth. What the episode makes particularly clear is that one cannot simply announce the existence of a new scientific discipline. Much of its existence depends

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

on recognition of practitioners outside the discipline, which is no trivial matter.

Chapter 7 begins by surveying the twentieth-century disciplinary history of paleontology, leading us to an investigation of what it meant to practitioners for evolutionary paleontology to be an autonomous discipline. Their understanding of paleontology's ability to contribute to evolutionary science rested on a range of philosophical commitments, from how to understand the practice of science to whether computer simulations of stochastic processes have ontological significance. Having properly framed their abstract disciplinary aims, chapter 8 examines three unsuccessful attempts to achieve them. Chapter 9 ends on a high note, showing how the long campaign to realize conditions (2) and (3) above resulted in desired effect: a new scientific discipline.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

7 The Philosophical Foundations of Disciplinary Autonomy

7.1 What Future for Historical Science?

> Paleontology is not a branch of either biology or geology, it is paleontology.
> —Norman D. Newell and Edwin H. Colbert, "Paleontologist: Biologist or Geologist?" (1948, 266)

In 1963, Stephen Jay Gould entered Columbia University as a graduate student under the tutelage of Norman Newell, a professor of geology and Curator of Invertebrate Paleontology at the American Museum of Natural History. Since his arrival at Columbia in 1947, Newell had helped to initiate many of the lines of substantive and methodological thought that would later come to characterize the "paleobiological revolution." In Gould's memory, Newell's influence over the future of paleontology approached messianic proportions:

> When virtually all paleontologists were trained as geologists and had no biological knowledge beyond the basics of invertebrate morphology, Norman Newell saw, virtually alone, that the most exciting future direction in paleontology lay in its relationship to evolutionary theory and to biological thought in general. I think that only a few very old-fashioned paleontologists would deny today that this prediction has been fulfilled and that American invertebrate paleontology is now in its most exciting phase since the era immediately following Darwin's *Origin of Species*. With his early monographs, and his persistent encouragement of biological thinking, Norman Newell was the godfather of this movement. (Stephen Jay Gould to Niles Eldredge, March 9, 1978, American Museum of Natural History Department of Invertebrates archives, quoted in D. Sepkoski 2012, 55)

Historian of science David Sepkoski vividly describes how Newell devoted much of his professional energy to the reorientation of the discipline of paleontology away from its affiliation with stratigraphy and petroleum

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

geology and back toward its historical role as a biological and, in Darwin's hands, *evolutionary* discipline. At the level of research foci, this would require a shift from the "traditional" work of description and classification of fossils for geological purposes, to the "understanding and interpretation of fossils and their life environment" (Norman Newell, "Instruction in Paleobiology": American Museum of Natural History Department of Vertebrate Paleontology archives, 67, 21). For Newell, the proper understanding and interpretation of the fossil record would require "an *evolutionary* viewpoint" (Newell and Colbert 1948, 267).

Newell's vision for paleontology's new disciplinary orientation drew much of its inspiration from the contemporaneous work of his professional mentor and colleague, George Gaylord Simpson, who had published *Tempo and Mode in Evolution* (1944) just prior to Newell's arrival at the American Museum of Natural History. *Tempo and Mode* had shown that much of the data special to paleontology, principally patterns and shifts in the fossil record, could be given an evolutionary interpretation that made these phenomena consistent with "general evolutionary theory" as articulated by Fisher, Haldane, and Wright during the Modern Synthesis (Simpson 1945, 45). To vastly understate matters, this was an important achievement for paleontology. The genetic work undertaken by Fisher et al. had set the course of theoretical biology for the foreseeable future, and Simpson's book used that work (notably the work of Dobzhansky 1937) to demonstrate the importance of the fossil record as a *biological* phenomenon. By pointing to cases where Darwinian expectations were borne out by the fossil record—series of gradual morphological progression—he made a convincing case for the genuinely biological nature of fossil patterns.

Even though this consistency argument was critical for paleontology's inclusion in the Modern Synthesis (Gould 1980), Simpson's aim in *Tempo and Mode* was considerably more ambitious than this. Ultimately he would lobby for the necessity of paleontological data for the construction of a complete evolutionary theory. In his own, strong words,

> it is not only that the theory is incomplete until it is supplemented by the inclusion of paleontology in the synthesis, but also that it may be quite wrong, even on its lowest levels, as long as there are apparent inconsistencies between theory and what the fossil record suggests as fact. (Simpson 1945, 58)

With the biological bona fides of fossil patterns fully established, Simpson was able to insist that an adequate evolutionary theory would need to

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

accommodate all the biological facts reflected in the fossil record, not just those that indicated gradual directional trends in morphology. As long as paleontologists could confirm the authenticity of an apparent signal in the fossil record, a complete evolutionary theory would have to account for it.

The argument in *Tempo and Mode* emboldened the campaign to transform paleontology in two important ways. First, the credibility that *Tempo and Mode* gave to the idea of interpreting fossil patterns as evolutionary phenomena provided a solid foundation on which one could argue for paleontology's status as an evolutionary discipline and, consequently, the need for students in paleontology to be trained primarily in biology and evolutionary theory. And, indeed, Simpson's effort to establish the biological and evolutionary significance of the fossil record was paralleled by a contemporaneous current within paleontology that emphasized the disciplinary and pedagogical implications of that significance. J. Brookes Knight made precisely this argument in his presidential address to the Paleontological Society in 1946, "Paleontologist or Geologist?," an unequivocal call for training in paleontology to be predominately biological in character:

> The paleontologist is dealing with the remains of once living things that are governed by precisely the same fundamental biological laws and subject to the same biological processes as are the living things of today. Obviously, unless he has a broad understanding of the principles and laws of biology, and a deep understanding of the more relevant of them, he cannot truly understand fossil organisms. (Knight 1947, 283)

Knight's address created a minor stir in the paleontological community, and Newell was quick to register his strong approval for the "real service to paleontology and geology for calling attention to serious traditional deficiencies in the training of invertebrate paleontologists," the group most closely associated with geological applications of paleontological science (Newell and Colbert 1948, 264).

Second, *Tempo and Mode* suggested a strategy by which paleontologists could make independent and distinctively paleontological contributions to evolutionary theory. If, as Simpson had argued, a complete evolutionary theory would have to account for patterns in the fossil record, then paleontologists' theoretical ambitions could be advanced by finding those patterns and developing the theories that could account for them. If some fossil pattern could not be credibly interpreted as reflecting Darwinian processes, that pattern represented an opportunity for paleontologists to

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

propose general evolutionary mechanisms that operated independently of the "lowest taxonomic levels." As Simpson rather pointedly remarked:

> There may be some *a priori* likelihood that the factors of evolution involved in large-scale evolution, effecting major morphological chances, and extending over millions of years, will be the same as those involved in speciation, but this can by no means be taken as a firm assumption. It is not even likely that those factors will combine in the same way and produce identical patterns or modes in the larger picture as in the smaller. (Simpson 1945, 57)

In sum, it had become clear to a number of people in the field that the maintenance and growth of paleontology as a science would depend on the degree to which training and scholarship in paleontology were geared toward becoming an evolutionary discipline. At the level of pedagogy, this required supplying the student with the theoretical and methodological tools needed to detect evolution in the fossil record; as Knight had stressed, "unless he has a broad understanding of the principles and laws of biology . . . he cannot truly understand fossil organisms." In terms of research foci, the way forward was rather poorly defined. It clearly would not serve paleontology's disciplinary interests to go from playing handmaiden to geology to playing handmaiden to biology, and so it would not be sufficient to use the fossil record merely to confirm the general evolutionary theory established in the Modern Synthesis. For the discipline to achieve the sort of autonomy to which Newell et al. aspired, it would have to generate some theoretical demand that the paleontological perspective alone could satisfy. Thus, Simpson's point that paleontology had a unique role to play in the evolutionary interpretation of properties of higher taxonomic levels was more than a claim about what a complete evolutionary theory required, or about how paleontologists might profitably pursue their evolutionary theoretic ambitions. It was a template with which they could achieve genuine disciplinary autonomy.

Let us pause here for a moment to examine how these developments connect with earlier themes associated with Darwinian populations. First, observe that the actors in our story clearly understood that the survival of paleontology as an independent discipline would depend on whether it could successfully *isolate* itself from the disciplinary trajectories of both geology and biology; paleontology's autonomy required that it was able to do more than take its direction from the scientific demands of these

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

other disciplines. Second, they saw that the key ingredient for initiating successful isolation was the "control of some defensible resource" (Pagel and Mace 2004, 278). In the specifically scientific cultural context before us, this amounted (per chapter 5) to the identification of a set of scientific problems that paleontologists alone were equipped to metabolize.

These two components—increasing isolation from other populations, and the monopolization of some resource—are central to the fulfillment of the *intrinsicality* condition for Darwinian populations, because each of them contributes independently to realizing a situation in which differences in the success of scientific practices in paleontology depend on how well those practices contribute to goals that are important to most paleontologists, but not to most non-paleontologists. On the one hand, isolating paleontology from the demands of other disciplines lessens the degree to which paleontological practices contribute to goals that matter as much or more to practitioners outside of paleontology than to paleontologists themselves. On the other hand, identifying a set of significant problems for which paleontology is uniquely well suited increases the degree to which it is possible for paleontological practices to contribute to goals that matter more to paleontologists than to non-paleontologists. This is precisely the emphasis we would expect to see in the very early stages of developing into a Darwinian population, because, more than any other features, these two components are essential for a group of practices to comprise a unit that can in principle be subject to evolutionary change.

Having said that, it was not clear precisely which fossil phenomena might provide fuel for distinctively paleontological evolutionary theorizing. Simpson's own *Tempo and Mode* flirted briefly with the notion of treating the origin of higher-level taxonomic units as a special sort of phenomenon that did not yield readily to analysis in terms of gradual adaptive evolution. Instead, Simpson proposed that certain of these cases be treated as examples of what he called *quantum evolution*, a process that he "believed to be the dominant and most essential process in the origin of taxonomic units of relatively high rank, such as families, orders and classes. It is believed to include circumstances that explain the mystery that hovers over the origins of such major groups" (Simpson 1944, 207).

There was no question that the fossil record reflected the phenomenon to which Simpson drew attention, and he was quick to stress that our knowledge of its empirical reality was due solely to paleontological

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

evidence (Simpson 1944, 207). Simpson's original explanation of this phenomenon in *Tempo and Mode* invoked an essentially "Wrightian" mechanism, whereby a population previously at equilibrium around one adaptive peak becomes unstable and is rapidly propelled by random drift across a valley and then back up another adaptive peak through standard directional selection. What distinguished this sort of evolution from evolution in the phyletic mode (evolution along a lineage) was the period of disequilibrium that characterized the shift from one adaptive zone to another:

> In phyletic evolution equilibrium of the organism-environment system is continuous, or nearly so, although the point of equilibrium may and usually does shift. In quantum evolution equilibrium is lost, and a new equilibrium is reached. There is an interval between the two equilibria, the biological analogue of a quantum, in which the system is unstable and cannot long persist without either falling back to its previous state (rarely or never accomplished in fact), becoming extinct (the usual outcome), or shifting the whole distance to the new equilibrium (quantum evolution, strictly evolution). (Simpson 1944, 207)

By 1953, however, the special status that Simpson attributed to quantum evolution had faded, and he now regarded it as a mere "limiting case" of ordinary gradual adaptive phyletic evolution.

Simpson's reinterpretation of the sort of unique fossil pattern that characterized quantum evolution was a setback for the effort to carve out a disciplinary space for paleontology. True, it was an evolutionary phenomenon that only paleontology was able to authenticate, and so represented an important step toward disciplinary autonomy. But by treating quantum evolution as governed by the same evolutionary mechanism that was dominant at the lower taxonomic levels—gradual modification of a population through natural selection—it ceased to be a phenomenon that only paleontology could explain. Thus, Simpson had deprived paleontologists of a chance to make an independent contribution to the family of general evolutionary laws. Because there was nothing particularly special about quantum evolution, there need be nothing particularly special about the laws governing it. Nevertheless, in broad outline the way forward remained clear: find fossil phenomena that were of unambiguous evolutionary importance yet resistant to Darwinian gradualist explanation and formulate the evolutionary laws that could account for those phenomena.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

7.2 "Geology Is a Science"

By the time Stephen Jay Gould entered Columbia in 1963, the charge led by Newell to transform paleontology into a thoroughly evolutionary discipline had achieved considerable momentum and had resulted in the development of a battery of new tools, questions, and interests with which one could approach the fossil record. These included "the investigation of broad patterns in the fossil record, the development of quantitative approaches to fossil databases, the study of the evolutionary significance of mass extinctions, and the creation of the sub-discipline of paleoecology," each of which was vital to the envisioned transformation (D. Sepkoski 2012, 54). The first three in particular stand out in connection with the template for achieving disciplinary autonomy mentioned above. The relevance of the investigation of fossil patterns is obvious, and quantitative approaches to database analysis were an attempt to uncover statistical patterns that themselves suggested the existence of real signals coming from the fossil record.

Mass extinction was a phenomenon to which paleontology had proprietary access, and Newell had been increasingly vocal in drawing attention to its existence and importance following his 1952 paper "Periodicity in Invertebrate Evolution." Here Newell pointed to some fossil phenomena of potentially major evolutionary significance. One was the regular emergence of new higher-level taxonomic groups immediately following a mass extinction among invertebrates, which he attributed to "a sort of adaptational vacuum" that provided ecological space for new kinds of occupants (Newell 1952, 384). The second was the related acceleration in evolutionary rates (measured as rates of genera diversification) that mass extinctions appeared to usher in (385). Newell had, in effect, attempted to provide one kind of explanation for the phenomenon of *eruptive evolution*, a term coined by invertebrate paleontologist Preston Cloud a few years earlier that referred to a fossil "pattern and sequence" characterized by

> (1) relatively sudden appearance of marked variability; (2) probable availability and proximity of a variety of ecologic niches; and, finally, (3) increased selective pressure, resulting in the weeding out of inadaptive or poorly adaptive radicles and leading to a more regularly channeled evolutionary phase for the particular stock involved. (Cloud 1948, 346)

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Mass extinction was one (but of course not the only) mechanism that could explain (2)—the appearance of available niche space—and could thereby explain the associated nonuniformity in morphology and evolutionary tempo across time.

Both of these themes—paleontology's access to special evolutionary phenomena, and paleontology's use of laws of nature to explain those phenomena—would come together in Gould's first published essay. Shortly after beginning his graduate training with Newell, Gould wrote a short article titled "Is Uniformitarianism Necessary?" that addressed a particular conceptual confusion which had plagued the literature in geology for over a century. Published in 1965 in the *American Journal of Science*, the article pointed to two uses of the term "uniformitarianism" in the geological literature—one substantive, one methodological—that had often been pitted against each other in a senseless cacophony that simultaneously undermined scientific progress within geology and the image of geology within the sciences generally. *Substantive uniformitarianism*, the empirical claim that the rates of geological change had been constant throughout history, "has not withstood the test of new data and can no longer be maintained in any strict manner" (Gould 1965, 226), a development that Gould attributed to the growing research program on mass extinction led by Newell and German paleontologist Otto Schindewolf. In contrast, *methodological uniformitarianism*, the assumption that the laws of nature remain constant across time and space, "remains vital to geologic inquiry" (226).

The importance of methodological uniformitarianism for geology has its roots in Lyell, from whom Gould approvingly quotes the following: "Our estimate indeed, of the value of all geological evidence, and the interest derived from the investigation of the earth's history, must depend entirely on the degree of confidence which we feel in regard to the permanency of the laws of nature" (Lyell, 1830, 1:165). Gould goes on to provide interpretation and defense, in what would come to be the first of many a dramatic flair:

> Once accepted, this uniformity ended the dichotomy between a contemporary world operating under constant and verifiable natural laws and a past incapable of purely scientific explanation. The entire geologic record, with all its evidence of vast upheaval and mass extinction, was, for the first time, integrated within the sphere of empirical investigation. (Gould 1965, 224)

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

What the assumption of nomic spatiotemporal invariance provides—what makes the scientifically impossible possible—is the warrant for projecting observed regularities to unexamined cases. But, argues Gould, there is nothing special about geology's commitment to methodological uniformitarianism. It is, after all, nothing more than "a statement of proper scientific procedure in general, independent of any particular substantive theory" (1965, 224):

> The assumption of spatial and temporal invariance of natural laws is by no means unique to geology since it amounts to a warrant for inductive inference which, as Bacon showed nearly four hundred years ago, is the basic mode of reasoning in empirical science. Without assuming this spatial and temporal invariance, we have no basis for extrapolating from the known to the unknown and, therefore, no way of reaching general conclusions from a finite number of observations. (Gould, 1965, 226)

Another feature of methodological uniformitarianism, one that requires but is not grounded in the assumption of nomic spatiotemporal invariance, is the notion that "reference need only be made to presently-observable causes in explaining past changes" (Gould 1965, 227). This is of course the famous *vera causae* principle on which Lyell had so heavily insisted. Thus, methodological uniformitarianism turns out to be the claim that all and only the known laws of nature are projectible. In Gould's words,

> methodological uniformitarianism amounts to an affirmation of induction and simplicity. But since these principles belong to the modern definition of empirical science in general, uniformitarianism is subsumed in the simple statement: "geology is a science." But since we consider geology a science, this affirmation has already been made by definition. Saying it again is at best superfluous and at worst confusing since it leads to the inference that our science has a powerful and unique guiding principle all its own. The unity of procedural assumptions, which binds the empirical sciences together, should not be obscured by terminology specific to one discipline. (1965, 227)

For Gould (at least, in 1965), to be a science is to use laws of nature in order to project from examined cases to unexamined cases. In this respect, geology is just like any other science.

There are two features of Gould's "Is Uniformitarianism Necessary?" that should be emphasized in connection with paleontology's strategy for achieving disciplinary autonomy outlined above. The first is the broader evolutionary—and associated disciplinary—implications of substantive

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

uniformitarianism. Substantive uniformitarianism, which asserts the uniformity of conditions or rates of geologic change, provided critical support for the notion that evolutionary change was slow and gradual across geologic time. As long as substantive uniformitarianism was assumed, leaps across morphospace represented in the fossil record could be attributed to poor preservation of evolutionary series; if substantive uniformitarianism is true and implies uniform evolutionary rates, and the fossil record suggests otherwise, the fossil record must be incomplete. Since all fossil patterns fit gradualist expectations—either in manifest physical detail or by inference—there would be no need for alternatives to the general Darwinian treatment. But once substantive uniformitarianism can no longer be maintained, as Gould thought it could not (owing to the work of Newell and Schindewolf), it becomes possible to treat leaps across morphospace (as well as less dramatic accelerations and decelerations in morphological divergence) as real evolutionary phenomena. In this light, the significance of Gould's argument can be appreciated as more than an attempt to clear up an annoying homophony. It is fundamentally an attack on the chief geological presupposition underlying the paleontological argument for phyletic gradualism (though of course it was not an attack on the argument for Darwinian gradualism generally; that had been strongly supported through the mathematical and empirical investigations of the Modern Synthesis). With the hegemony of phyletic gradualism in question, and with the principal evidence for skepticism coming from the fossil record, paleontology becomes uniquely qualified to document and explain non-Darwinian alternatives.

The second feature worth discussing is Gould's interpretation of the phrase "geology is a science." Conceiving of science as projection to unobserved cases on the basis of spatiotemporally invariant laws implies certain things about how scientific work is to be undertaken. Since the paleontologist's scientific goal is to understand unobserved (because ancient) life, and since the only way to learn about unobserved cases is to use our knowledge of observed cases to make inferences about them, and since the only inferential mechanism capable of underwriting the application of present knowledge to unobserved cases is the assumption of nomic spatiotemporal invariance, it follows that the only way for the paleontologist to achieve his scientific goal is to reason on the basis of laws of nature. Now, some fossil phenomena (e.g., finely graduated, continuous morphological series) may

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

be subsumable under known laws (e.g., Darwinian gradualism). In those cases, argues Gould, the methodological precept of simplicity requires that we appeal to those laws known to be operative. However, when fossil phenomena like leaps across morphospace resist subsumption under known laws, the paleontologist's ability to do science depends on his ability to discover new laws of nature that can account for the recalcitrant fossil evidence. For Gould, then, fossil phenomena not subsumable under Darwinian gradualism translate into an opportunity—indeed, a mandate implied by the very definition of science itself—to discover the laws governing those phenomena.

On the basis of the historical evidence known to me, it is not discernible whether Gould inherited these views from Newell or came to Columbia with a philosophy of science and implied disciplinary agenda already in place.[1] But it is also not relevant. It is sufficient for this stage of my argument to have shown that Gould was not only attracted to paleontology's new direction but was moved to provide a general argument for its propriety by showing that new direction to be a consequence of the intrinsic character of science—the application of, and thus the search for, laws of nature. In essence, this was Gould's metastrategy for developing a distinctively paleontological Darwinian population.

7.3 The Conceptual Foundations of Paleontological Laws

> There are nomothetic undertones to the results of evolution—the principle of natural selection is among them—and it is here that our laws must be formulated. They must be based on immanent processes that produce events, not on the events themselves.
>
> —Stephen Jay Gould, "Dollo on Dollo's Law" (1970a, 209)

One of Stephen Jay Gould's many contributions to historical scholarship was his study of Dollo's Law, named for Swiss naturalist Louis Dollo (1857–1931). For us, the real value of this study lies not in what it tells us about Dollo, but in what it tells us about Gould, for it is here that all of the notable streams of Gould's scientific thought come together; laws, contingency, anti-inductivism, hypothesis testing—each of them woven into an impressively coherent narrative that sheds significant light on Gould's philosophy of science, his scientific method, and the "paradoxical situation" to which he alludes in the final pages of *The Structure of Evolutionary Theory*.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Titled "Dollo on Dollo's Law: Irreversibility and the Status of Evolutionary Laws," the purpose of the essay is to correct various misinterpretations of Dollo's claims concerning the irreversibility of evolution and then bring the correct interpretation to bear on the study of evolution. Dollo claimed that physical structures do not evolve more than once, a fact that he (channeled through Gould) attributed to the physical complexity of organisms. Given that each specific morphological change in a lineage is a very, very low-probability event, the probability of the same structure evolving twice on a lineage is negligible. Thus, Dollo's Law: morphological changes along a lineage are irreversible, and so "a whole organism never reverts completely to a prior phylogenetic stage" (Gould 1970a, 201). Here Gould argues that Dollo's Law of the irreversibility of evolution was not an empirical generalization derived from the fossil record but rather a special case of the principle that "When a phenomenon reaches a sufficient degree of complexity, requiring a sufficient number of independent steps for its realization, repetition becomes 'absolutely unimaginable—there are too many other possibilities, the probability is nil'" (199).

Readers familiar with contingency-Gould will instantly recognize what Gould found important about Dollo's Law: it serves as the conceptual basis for his views about contingency of evolutionary history. Gould's favorite metaphor for the history of life was that of a "tape," which, if erased, rewound, and replayed, would thrust life in an entirely new direction, resulting in a different panoply of phenotypes. It is Dollo's Law—the principle that Markov chains of complex events are not repeatable—that supports the evolutionary intuitions evoked by "replaying the tape of life."

But Gould saw the significance of Dollo's Law as extending beyond its biological meaning. As he observed, "Dollo's law is not an adjunct of evolutionary theory. It is a statement, framed in terms of animals and their evolution, of the nature of history; or, put another way, it is an affirmation of the historical nature of evolutionary events" (Gould 1970a, 208).

For Gould, this affirmation held implications for "the entire enterprise of lawmaking for phylogenetic results" and, consequently, for how a science of paleontology ought to look (Gould 1970a, 209). Fossils record historical events, and history does not repeat itself. Each fossil records a highly improbable sequence of events that could have been otherwise—indeed, would have been otherwise—were we to "replay the tape of life." When we attempt to formulate laws on the basis of the physical features of particular

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

fossils, we succeed only in describing what a bunch of historical accidents have in common. But there is nothing lawlike about the highly contingent phenotypic states recorded by the fossil record since, ex hypothesi, they could (would) have been otherwise. In effect, the truth of Dollo's Law implies that an "enterprise of lawmaking for phylogenetic results" is doomed to fail because laws must describe states that could not have been otherwise, and Dollo's Law implies that phenotypic states would have been otherwise. Here is Gould at his best:

> Simpson has distinguished immanent from configurational properties of the universe (the former as "the unchanging properties of matter and energy and the likewise unchanging processes and principles arising therefrom"; the latter as "the actual state of the universe or of any part of it at a given time). Laws are framed for immanent properties: we are not interested in the melting behavior of a particular ice cube but in the properties of water in general. Physics rarely deals with the configurational; if its formal structure is lawlike, this is because it has excluded the configurational from its domain. The error made by reductionists who attempted to formulate laws for the results of evolution was that they assumed a similar focus for biology and physics. But biology often deals with the configurational and the search for so-called historical laws among such properties is not a fruitful endeavor. (Gould 1970a, 209)

The difference between physical science and biological science, as he saw it, was primarily a matter of "difference in emphasis," and the emphasis within paleontology on "the configurational"—contingent evolutionary outcomes—had inhibited our ability to discover laws.

But if the laws are not to be found among the fossils, then where? His prescribed remedy was for our lawmaking efforts to be "based on the immanent processes that produce events, not on the events themselves" (Gould 1970a, 209). And how are we supposed to do that, exactly? The emphasis on "the configurational" was, for Gould, a symptom of paleontology's addiction to "inductivism," a refusal to venture outside "the observational mode of traditional natural history" (Gould 1970b, 88). Consequently, he saw the success of lawmaking endeavors in paleontology as hinging on a shift away from an observational style of inquiry—describing fossils and the like—to an approach that focused on constructing models of evolutionary processes and using fossils to test those models. I chronicle his forays beyond "the observational mode" in chapter 8.

There is thus an impressively straight line from Gould's embrace of the contingency of evolutionary results to his embrace of modeling as the most

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

fruitful approach to paleontological inquiry. But while the significance of contingency for Gould lay, at least initially, in what it implied for how we might fruitfully study natural history, its most famous role was as a premise in his long argument against adaptationism (Beatty 2006). Paradoxically, his first attempt to design a new paleontology, employing the modeling approach required by evolutionary contingency, produced a research program so spectacularly and shamelessly adaptationist that it makes the Sociobiological targets of Gould and Lewontin (1979) look mild by comparison. One can hardly believe he would emerge—and be remembered—as adaptationism's most vocal critic.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

8 Attempted Paleontology

> Convicted of a crime I didn't even commit. Hah! Attempted murder? Now honestly, what is that? Do they give a Nobel Prize for attempted chemistry? Do they?
> —Sideshow Bob from *The Simpsons*

No, they do not. It is one thing to affirm the scientific importance of the search for laws of nature, quite another thing to conduct that search in a way that appears credible, potentially fruitful, and capable of underwriting claims of disciplinary autonomy. In this chapter I document three of Gould's attempts to build a laws-oriented research program that could act as a stepping stool to paleontology's seat at the high table of evolutionary theory. All of these attempts failed. More importantly, they all failed for reasons that are straightforwardly predicted by the model of disciplinary emergence under consideration. That model states that the emergence of an autonomous discipline requires the identification of a set of problems whose solution is in high demand and proprietary access to the perspective best suited to solving those problems. If either of these conditions is missing, the fledgling discipline will fail to take flight; there can be no Darwinian population without them. Each of the failures described below, I will argue, is directly attributable to the absence of one of these factors.

Although none of the campaigns I discuss here would ultimately serve to deliver paleontology to the promised land, each of them shares an important feature that would prove critical in the eventual transformation of paleontology into a distinct branch of the evolutionary sciences. This is the element of model testing to which Gould began drawing attention in the late 1960s. In the background of each laws-oriented research program was an unshakeable faith that modeling was the way forward, and that success

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

was merely a matter of modeling a fossil phenomenon distinctive and important enough to attract the attention of the evolutionary community.

8.1 First Attempt: The Science of Form

> Our ultimate goal in the study of a phyletic lineage is the explanation of each morphological change in terms of its selective advantage.
>
> —Stephen Jay Gould, "Allometry and Size in Ontogeny and Phylogeny" (1966, 621)

> I fledged in science as a firm adaptationist, utterly beguiled by the absolutist beauty (no doubt, my own simplistic reading of a more subtle, albeit truly hardened, Modern Synthesis) of asserting . . . that all aspects of organismal phenotypes, even the most trivial nuances, could be fully explained as adaptations built by natural selection.
>
> —Stephen Jay Gould, *The Structure of Evolutionary Theory* (2002, 40)

Gould's earliest work was in allometry, publishing a coauthored paper on the biological meaning and use of the power function, $y = bx^k$ (often called the "allometric equation") in 1965 and an important essay review of the subject in 1966 (White and Gould 1965; Gould 1966). The immediate focus of this work is, predictably, the coordinated changes in size, shape, area, and volume that take place along lineages. But his broader concerns are with the evolutionary explanation for those changes (which he believed resided squarely with the causal power of natural selection) and fruitful methodologies for studying natural history.

The ubiquitous correlation between changes in size and changes in shape along lineages suggest that certain kinds of morphological change must necessarily accompany others. Thinking of organisms as kinds of physical structures provides a simple and general explanation as to why: "certain shape alternations are mechanically required by size increase," independently of whether those structures are biological or not (Gould 1966, 588). But the fact that biological lineages are, all on their own, responsive to those mechanical requirements requires special explanation. The general account presented by Gould was that "constant area-to-volume ratios" are "an adaptive necessity for many organic relationships," a fact that generates strong selective pressure within a lineage to maintain those ratios (588). When organisms undergo an increase in size, their area is increasing

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

as the square of linear dimensions, but their volume is increasing as the cube of linear dimensions. For example, a little elephant that doubled in size would see a corresponding fourfold increase in surface area, and a corresponding eightfold increase in volume. For functions such as respiration that are closely tied to surface area, the extra physiological demand generated by an eightfold increase in volume will overwhelm the comparatively small increase in functional capacity provided by the fourfold increase in surface area. Thus natural selection will put pressure on a lineage undergoing size change to find ways of accommodating the new demands placed on it by its volume increase, and that pressure will often result in changes in shape that preserve the lineage's prevailing adaptive ratio of surface area to volume.

By 1970, Gould had come to view the approach to allometry described in his 1966 essay review as the instantiation of a more general mode of paleontological inquiry, one that had all the markings of a new autonomous scientific discipline (Gould 1970b). Dubbed "the science of form," the hallmark of the new discipline was the mechanically oriented study of adaptation outlined by then-paleontologist (turned historian of science) Martin Rudwick in a series of papers in the early sixties (Rudwick 1964) and impressively executed by a handful of researchers, notably the young David M. Raup (1966). The essence of the mechanically oriented approach was to first adopt as a methodological assumption the idea that the relevant morphological feature was an adaptation designed by natural selection to perform a specific function. But performing that function requires that certain design criteria be met, and natural selection will favor those forms that most completely satisfy those criteria. By developing an engineering-style optimality model of what it would take to meet the design criteria required for performing that function, we can derive quantitative expectations for what a character should be like if it is indeed an adaptation designed to perform that function. We then compare those quantitative expectations to actual fossils to see whether the fit between prediction and observation is adequate.

Measured against either his particular scholarly interests or his disciplinary ambitions, it is no mystery why this style of inquiry resonated with the young Gould. At the scholarly level, there is a strong resemblance—arguably a unity—between Gould's own work, which employed the allometric equation as a model of a selective constraint on organismal form, and Rudwick's engineering analysis, which used mechanically derived optimality models as selective constraints. One suspects, however, that the

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

disciplinary implications of the approach were where Gould's most intense interest resided. Most importantly, the new approach fell squarely outside the inductivist, observational mode for which he largely blamed paleontology's failure to mature into one of the "grown-ups" of evolutionary theory. Rudwick seems to have agreed (at least, in part), asserting that "if our primary purpose is the construction of phylogenies, or taxonomic classification, we shall be blind to the clues that might lead to successful functional inferences" (Rudwick 1964, 35). From the perspective of the new science of form, the morphological similarities, whose diligent recording had been the dominant research focus of paleontologists, were mere "clues" that could only serve as suggestive of certain functions, the discovery of which being the key to understanding the evolutionary basis of morphological change.

Apart from its eschewal of narrow inductivism, the specific shape of the new science of form embodied Gould's vision of a fruitful evolutionary lawmaking enterprise. For one thing, the methodological precepts of the new science had inquiry beginning at the appropriate level; as he describes the approach, "Idealized models are favored over actual specimens" (Gould 1970b, 77). By constructing models of evolutionary processes—specifically, the optimizing process of adaptation—paleontological inquiry could focus on the properties of adaptation in general, rather than on the morphological change of a particular taxon (a paleontological variation on Simpson's distinction between immanent and configurational properties mentioned above). This shift toward the modeling of "the immanent *processes* that produce events, not on the events themselves" would allow paleontologists to study the "nomothetic undertones" of the fossil record directly, rather than embarking on the fruitless search for "'laws' . . . based . . . on configurational properties," which were in Gould's view nothing more than "descriptive generalizations of low-order probability that describe some common regularities without explaining anything" (1970a, 209).

But the value of modeling was not limited to its nomic fecundity; the largess to other dimensions of scientific inquiry were also important. Gould had remarked in his 1966 essay review on the fact that the power function $y=bx^k$ was used far more widely ("almost exclusive" were his words) than anything else for capturing fossil data, which was significant because other formulas achieve a statistical fit that is consistently superior to the power function in individual cases. If other formulas were more quantitatively accurate, why was the power function the dominant modeling strategy?

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Gould's explanation for its popularity turns out to be closely related to his critique of the focus on configurational properties:

> Closer fits can always be obtained by adding further terms to a power series $y=a_0+a_1+a_2x^2 \ldots$ but this method of data abstraction offers little hope of increasing biological understanding. The possibility of interpreting parameters is here sacrificed for details of fit applying only to the particular case under consideration—detail which may, moreover, be only the random effects of small sample sizes. Adequate statistical fit, while certainly the most important consideration in choosing a mathematical expression for data trends, is not the only attribute of a "good" formula. . . . Sholl (1954, 225) emphasized the importance of choosing "a reasonably simply functional expression involving the minimum of non-interpretable parameters." It is for all these reasons—adequate statistical fit in a great number of cases, simplicity, and interpretability—that the power function $y=bx^{\alpha}$ has enjoyed almost exclusive use in allometric studies. (Gould 1966, 596)

The message here is that a more accurate description is but one of many scientific goals, and a relatively minor one at that. The trouble with the focus on descriptive generalizations—even good ones, like Cope's Law or Williston's Law—was that they did not "explain anything" (Gould 1970a, 209). We scarcely do better by buying greater descriptive accuracy for a particular fossil series in the form of uninterpreted parameters, at the price of dramatically reducing a formula's range of applicability and our ability to understand and explain morphological trends. In both cases, the focus on "the configurational" "is not a fruitful endeavor" (209).

The model-based "quantifunctional approach" Gould advocated, however, *was* "a fruitful strategy" (Gould 1970b, 78), the power function in particular enjoying "a long history of fruitful use" (Gould 1966, 595). Beyond its superior talent for uncovering laws, the model-based approach allows us to understand the mechanical or biological underpinnings of those laws, which in turn fosters our ability to extend models to other systems. By designing a model in a way that reflects our beliefs about how nature works, we entitle ourselves to the inference that our beliefs about nature contain some element of genuine understanding—when there is an adequate match between model and observations. In this way, models also stand poised to increase our understanding of other domains of phenomena. Where we have reason to believe that multiple domains are subject to similar kinds of influence, a model's ability to fit data in one domain suggests that it might fruitfully be applied to, and thus offer similar insight into, other domains. As Gould would remind his fellow paleontologists over the next few decades,

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

the accommodating practice of descriptive generalization offers no such insight:

> Raup (1966) recognizes two potential explanations for the restriction of molluscan order to certain regions of [morphospace], (*1*) phylogenetic accident (insufficient time or chance failure to populate the entire [morphospace]) or (*2*) functional, usually mechanical, necessity. Raup claims that "to draw this former conclusion is to disregard the question and thereby to ignore the possibility of a rigorous functional explanation" (1966, p. 1190). Testing the functional hypothesis is the best approach because its rival cannot be proven. In trying to apply the functional postulate, we might be stymied, but this would be the strongest potential support available for the idea of phylogenetic accident. In fact, we have not been stymied, rather stimulated. (Gould 1970b, 96)

In his view, optimality models were the ideal approach to understanding the mechanisms responsible for morphological trends, since a positive fit suggests adaptation to a specific function, and a lack of fit suggests lack of adaptation to that function. The more functions we test for, and the more lacks of fit we observe, the stronger the support is for the claim that these morphological trends are not caused by natural selection, that they are phylogenetic accidents. But unless we test for function, we won't have much support for any causal claims.

8.2 Science Delayed

> We celebrate, in this symposium, a quiet renaissance of biology's oldest subject. During the past decade, like a thief in the night, morphology has surreptitiously become interesting again. The descriptive concerns of cataloguers, constrained by their jargon to communication with a few other specialists, are yielding, or at least accommodating, to analytical and explanatory approaches . . . I believe, in short, that we are approaching a true "science of form"—a claim that has not been made with justification since Cuvier's time.
>
> —Stephen Jay Gould, "The Shape of Things to Come" (1973, 401)

> Finally, if I could, in a species of Devil's bargain, wipe any of my publications off the face of the earth and out of all memory, I would gladly nominate my unfortunately rather popular review article on "Evolutionary paleontology and the science of form" (Gould, 1970b)—a ringing paean to selectionist absolutism, buttressed by the literary barbarism that a "quantifunctional" paleontology, combining the best of biometric and mechanical analyses, could prove panadaptationism even for fossils that could not be run through the hoops of actual experiments.
>
> —Stephen Jay Gould, *The Structure of Evolutionary Theory* (2002, 41)

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Gould would continue witnessing to the new science of form through 1973, when he hailed its second coming while reporting on a symposium that he convened at the first meeting of the International Congress of Systematic and Evolutionary Biology for the purpose of showcasing the burgeoning movement. The following year he would publish a study of the relationship between body size and antler size in the "Irish Elk" *Megaloceros Giganteus*, arguing that they were allometrically correlated (Gould 1974). Although the article did not contain the proselyte prose of earlier pieces, he continued to preach the adaptationist gospel, concluding in line with previous studies that the correlation between body size and antler size was maintained by selection.

What is distinctive about this article, however, is the replacement of the kinds of quantitative optimality models of evolved function that characterized his earlier work with just-so storytelling, a crime for which Gould would soon be anointed Supreme Executioner during the Sociobiology Wars. Whereas the standard adaptationist explanation for the body size–antler size correlation had been direct selection for increased body size and allometrically mandated increase in antler size, Gould argued that this view rested on an unproven assumption—namely, that Irish Elk antlers were "weapons . . . developed beyond their optimum" (Gould 1974, 212). It's certainly true that if the antlers are weapons, then they are well beyond their optimum size: "90 point antlers, mounted with tines pointing backward on a 5 pound skull cannot be regarded as well-designed for such purpose" (217). He took this to imply that the antlers could not have been under direct selection for increased size. But Gould saw no reason to accept the assumption that Irish Elk antlers were weapons, and thus could not be persuaded to accept any conclusion other than that "the immense antlers of *Megaloceros* were advantageous in themselves." The more plausible hypothesis, he suggested, was that the antlers' incredible size was an adaptation produced by sexual selection, the adaptive function being "to establish dominance and win access to females by display and ritualized combat" (217). Consequently, increased body size may have been directly selected for independent reasons, or indirectly selected because of its allometric association with antler size. To be fair, Gould's supporting argument for the sexual selection hypothesis is far stronger than the speculative revelry for which sociobiology would later achieve notoriety. But it lacked the chief strengths of the original model-based approach, around which the

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

science of form was allegedly being organized: mathematical models and the quantitative predictions they entailed.

The comparative lack of rigor exhibited by his antler adaptationism may have presaged Gould's eventual disaffection with the "quantifunctional" optimality modeling approach to morphology because the Elk argument stumbles on precisely the weaknesses for which he would eventually blame the approach's general failure—namely, "Rudwick's analysis of morphology by comparison with ideal 'paradigms' or engineering optima (1961). It foundered," Gould proposed, "on the difficulty of specifying function a priori (in order to know what paradigm [optimality model] to construct)" (Gould 1980, 102). As he suggests here, whether these analyses are probative depends on whether we have an optimality model against which to compare our data. Viewed by those lights, what makes the antler adaptation argument weak by comparison with his case for the selective influence of the power function $y = bx^k$ (and more like a sociobiological just-so story) is the difficulty of constructing—and thus the absence of—a quantitative optimality model that can tell us what should be quantitatively true of antler size if they are designed "to establish dominance and win access to females by display and ritualized combat" (Gould 1974, 217).

But the optimality model-based approach to morphology had not been a total wash. Despite its failures, it "also generated fruitful models and hypotheses," having "clarified procedure and directed attention to the central problem of non-optimality" (Gould 1980, 102). Why were these two effects—procedural clarification and disclosure of the problem of non-optimality—so significant for Gould? The claim that Rudwick's optimality analysis "clarified procedure" can be read as saying that it clarified what it takes to demonstrate adaptive function, and on that basis I would argue that the comment is directed toward the sociobiological conflagration that was raging at the time; it's no accident that his famous coauthored "Spandrels" paper was published the previous year. Because support for the sociobiological adaptationist hypotheses was one of the central issues during the Sociobiology Wars (as he pointed out in his first critical essay in 1978), it is easy to appreciate Gould's sentiment that an improved understanding of what it takes to study adaptive function was an important development.

What of Gould's assertion that the optimality model–based approach was "fruitful" in that it "directed attention to the central problem of non-optimality"? (Gould 1980, 102). After 1973, Gould would never again attempt to rally paleontologists in the service of the science of form, which

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

he had worked tirelessly to develop in a variety of different dimensions—empirically, with his own work on allometry; methodologically, with his advocacy of Rudwick's approach to functional analysis; and professionally, through the organizing of symposia and other mechanisms of disciplinary cohesion. By 1980, he would pronounce the movement effectively dead:

> The flowering of functional morphology has yielded a panoply of elegant individual examples. . . . As particular studies of eminent utility in their own domain, these efforts have had great value—for we do need to know what this animal was doing in that place. But I think that higher hopes were once held (Rudwick 1964). (Gould 1980, 101)

Based on Gould's now-dim assessment of the prospects for optimality modeling, one might reasonably infer that he would consequently give up on the science of form. After all, his expressed hope and conviction had been that the science of form would coalesce around the "fruitful" lawmaking strategy of mathematical modeling, thereby revealing the new evolutionary laws that would bring paleontology to the high table of evolutionary theory. Once the science of form loses its modeling powers, its ability to discover laws of nature would seem, from Gould's point of view, to disappear.

Although this line of reasoning is intuitive, it turns out not to explain why Gould dissolved the science of form before it had fully emerged as an autonomous discipline. The real issue, I argue, was the failure of the science of form to deliver on one of the criteria for achieving disciplinary autonomy; in this case, a failure to identify a set of problems whose solution was in sufficiently high demand. As Gould himself remarked,

> I strongly suspect that such work, though intrinsically valuable for its elucidation of cases, will furnish no new or expansive generalizations . . . I, at least, once harbored the naive belief that a simple enumeration of more and more cases would yield new principles for the study of form (Gould 1970b). *But . . . who doubts that animals tend to be well designed?* (Gould 1980, 102; emphasis added)

The practical difficulties associated with employing optimality analysis were not the death knell for the science of form. Rather, the science of form would have fizzled out even if rigorous optimality analysis had been easy. This is due to the fact, here acknowledged by Gould, that the science of form was strictly in the business of demonstrating the influence of an evolutionary mechanism (natural selection) whose central importance was already accepted by everyone. Thus, Gould's science of form was never a good candidate for catalyzing the emergence of paleontology as an

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

autonomous evolutionary discipline, being as it was dedicated to investigating a set of problems whose solution was not at all in demand:

> *My own increasingly ambiguous feelings* about standard functional morphology *arise from its failure to generate new theory and its adherence to a conventional attitude that has, in my opinion, been both mined out and over-extended*. Too much of the work, while elegant in its application and eminently useful in the analysis of individual taxa, has not gone beyond the Newtonian premise that organisms work well. (Gould 1980, 111; emphasis added)

This brings us to the second reason for Gould's appreciation of the enlightening failure of optimality modeling. Because Darwin had already explained adaptation with the principle of natural selection, paleontologists would have to look elsewhere for a scientific problem that demanded solution with sufficient force to propel them to disciplinary autonomy. True, Darwin had disclosed the general evolutionary principle that explains why organisms tend to be well designed, but he had not been able to offer a general explanation (in the form of a distinct evolutionary mechanism) for why organisms tend to not be optimally designed. Nor could he have, which was essentially Gould's point. Just as only optimality modeling could rigorously demonstrate design, so too was it uniquely capable of demonstrating deviation from design (this, of course, would be in the lucky cases were function was specifiable a priori). This was a lesson he had learned from Raup (1966), whom he quoted to that effect in 1970, when his (Gould's) adaptationist science of form was building steam (Gould 1970b: 96, quoted above). It was only 10 years later, while attempting to salvage something from the smoldering science of form, that he would come to appreciate that same lesson's true significance:

> Interesting answers demand new questions, and empirical studies cast explicitly in their light. I suspect that functional morphology will fulfill its promise when it probes the situations in which animals are *not* well designed—developmental, phyletic, and architectural constraints as marks of history (Seilacher 1972; Gould and Lewontin 1979). As a key issue, I suggest: why is morphological space so sparsely populated, but so clumped where it is occupied? How much of clumping and non-occupation reflects good and untenable design (Raup 1966), how much the constraints of genealogy? (Gould 1980, 101)

From the ashes of an adaptationist science of form would, he hoped, emerge a decidedly anti-adaptationist version, but not because we've discovered that organisms tend to not be well designed, or because we can't routinely and effortlessly disclose adaptive function. Rather, because there are many

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

ways for them to be well designed, and yet the fossil record indicates that their actual design options appear to be highly constrained. This pervasive phenomenon is of very general significance, and its explanation might very well be the sort of new and general evolutionary law that the adaptationist science of form had attempted to disclose. Unlike its predecessor, the new science of form would identify the kind of "new questions" that "interesting answers demand"—a "central problem" whose solution demands the attention general evolutionary scientific community and fosters the fledgling phoenix in its effort to land a seat at the high table.

8.3 Second Attempt: "An Alternative to Phyletic Gradualism"

> This body of thought might serve as a basis for a new and general theory of macroevolution; this, obviously, is not a dispassionate opinion.
>
> —Stephen Jay Gould, "The Promise of Paleobiology as a Nomothetic, Evolutionary Discipline" (1980, 111)

Gould's campaign to forge the science of form was contemporaneous with what is perhaps his most well-known attempt to achieve disciplinary autonomy for paleontology: his advocacy, with Niles Eldredge, of *punctuated equilibrium*. Gould poured an enormous amount of time and passion into his effort to sell punctuated equilibrium to the evolutionary community, convinced that paleontologists had succeeded in identifying a special evolutionary phenomenon whose explanation would finally claim a space for them in the pantheon of evolutionary theory. In one respect, he was right: punctuated equilibrium was a special evolutionary phenomenon, and its authentication was the exclusive purview of paleontology. But punctuated equilibrium was neither new nor mysterious. From the perspective of our philosophical model of disciplinary emergence, this did not bode well for Gould's disciplinary pretensions. If punctuated equilibrium was already well understood, then its explanation was not in high demand. Thus, it could not serve as a basis on which to claim the autonomy of a new evolutionary discipline.

8.3.1 Eldredge's Insight

Niles Eldredge studied trilobite evolution at Columbia under Gould's advisor, Norman Newell. Part of his dissertation was published as an article in *Evolution* in 1971, in which he pointed out that

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> the concept of gradualism, an important aspect of geological thinking (see Simpson, 1970), has permeated paleontologic thought to the extent that all phylogenetic change is generally conceived to occur by small increments over vast periods of time. This dominantly *phyletic* model of transformation, stressing the importance of time and the aggregation of large numbers of small steps of morphological change, has underlain most paleontological discussions of the origin of new taxa, including species. In fact, this *phyletic* model applies not only to strict cases of phyletic transformation (i.e., linear trends in which arbitrary segments are viewed as 'new' taxa), but also to most discussions of divergence, where new branches in a phylogenetic tree are considered as gradually diverging stocks from a parent group. (Eldredge 1971, 156)

Eldredge is here drawing attention to the fact that phyletic gradualism no longer functions as an empirical hypothesis on which fossil evidence might be brought to bear, but is instead a lens through which all fossil data are filtered. Because phyletic gradualism is assumed in our encounters with the fossil record, both linear trends within a lineage (like increasing size, or successive elongation of certain structures) and the appearance of new lineages—that is, speciation—are both seen as gradual, step-by-step accumulation of morphological differences.

Eldredge's original paper went on to point out that, whatever merits the gradualist assumption may have had with respect to understanding phyletic transformations, it was radically out of step with the model of speciation that "has gained nearly total acceptance among current evolutionary biologists" (1971, 156). That model—the *allopatric model*—stated that speciation occurs primarily through the geographical isolation of a small sample of a wild population. In addition to being much smaller, the sample or "founder" population will tend not to be representative of the parent population with respect to how much genetic variation it contains. Both of these properties—reduced size and reduced variability—increase the probability that the alleles that go to fixation in the new, isolated population will be different from those of the parent population (reduced size increases the probability that a given allele's fitness and its actual frequency in future generations will diverge; reduced variability increases the probability that any one of the parent populations' alleles [including the fittest alleles] may simply not be in the geographically isolated sample). In turn, the increased probability that the parent and founder populations' fixed alleles will differ from each other, coupled with (1) the fact that the populations' members are no longer interbreeding and (2) the fact that adaptive requirements in

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

the two populations are likely to differ, suggests that the two populations will rapidly diverge genetically and, consequently, morphologically. This implication of the allopatric model contrasts with the slow and steady accumulation of minor differences between subgroups within a parent population, culminating in an insensible, anticlimactic splitting of the lineage.

The fact that the allopatric model was the predominant model of speciation among working biologists seemed sufficient, in Eldredge's opinion, that it

> be substituted in the minds of paleontologists for phyletic transformism as the dominant mechanism of the origin of new species in the fossil record, and that the allopatric model, rather than gradual morphological divergence, is the more correct view of the processes underlying cases of splitting already documented by numerous workers. (Eldredge 1971, 156–157)

But deference to the community of evolutionary biologists was not the only consideration involved in his case for why allopatric speciation was the "more correct view." Two important phenomena had emerged out of Eldredge's graduate work on the trilobite *Phacops rana*. One was that *Phacops* morphology was remarkably stable, typically exhibiting no change across millions of years; the other was that when the trilobites did change, it happened at an extraordinarily accelerated pace, taking "something between 5,000 and 50,000 years—a mere blink of an eye geologically speaking" (Eldredge 2008, 113).

For Eldredge, these phenomena were significant because they each in their own way defied the traditional expectation of gradual evolution. If evolution were a uniformly gradual process, we would expect to regularly see periods of slow, persistent, and insensibly graded directional change within populations in the fossil record—the process described by Darwin—entombed in the geological column. "On the contrary, the usual case, at least in Paleozoic epeiric [inland sea] sediments, is for the observer to document no change" (Eldredge 1971, 160). Relatedly, if evolution were a uniformly gradual process, we would also expect to not see "rapid burst[s] of change" of the kind observed during *Phacops* speciation events, and certainly not with the frequency there exhibited (Eldredge 2008, 112).

Figures 8.1B and 8.1C depict the processes that we would expect the ideal fossil record to reflect, using the gradualist model. Figure 8.1B, which represents the "standard phyletic model," was at the time of Eldredge's graduate work "probably the process most commonly cited to explain such

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

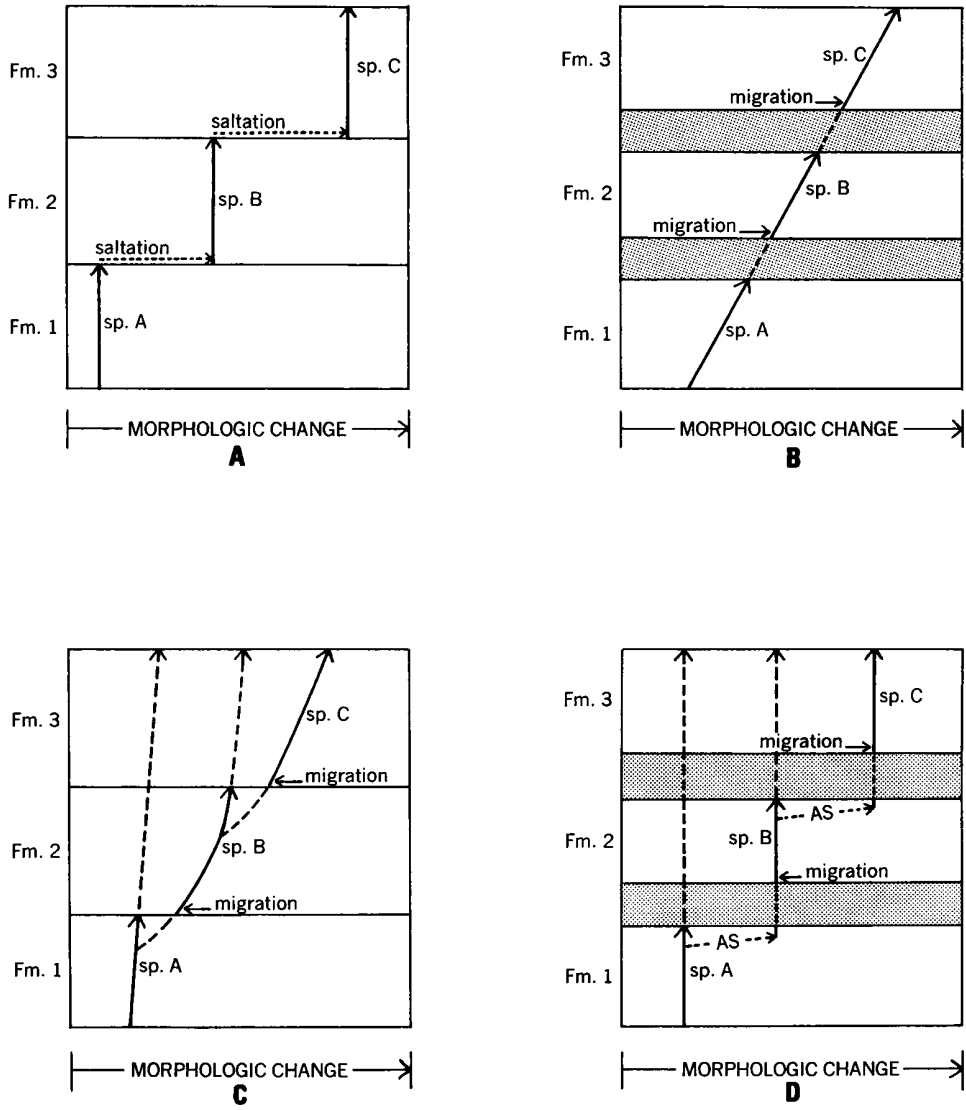

Figure 8.1
Four models of morphological change (from Eldredge 1971).

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

occurrences" as he observed in *Phacops* (Eldredge 1971, 160). Here we see slow and steady morphological change that is uniform both in rate and direction, taking place within a single lineage arbitrarily segmented to reflect current taxonomy. In this model, speciation is a purely phyletic process, where one species gradually transforms into another species without the occurrence of any branching events.

Figure 8.1C is another gradualist model, only with speciation depicted as a branching process rather than a purely phyletic one. Like figure 8.1B, most morphological change occurs on particular lineages, reflecting expectations associated with a phyletic gradualist understanding of evolution. In contrast to the process represented in figure 8.1B, whether something qualifies as a species depends on whether a branching event occurs, resulting in the creation of a distinct genetic lineage that (unlike the purely phyletic model) may or may not eventually cause the disappearance of its parent lineage. Figure 8.1D, the allopatric model, fits the rapid morphological evolution that Eldredge had observed in *P. rana*. Here, the degree of morphological difference characteristic of distinct species is achieved during the relatively short time in which branching actually occurs—precisely as was suggested by the *P. rana* fossil record.

Remember that Eldredge's trilobites also exhibited a rather surprising tendency to remain unchanged for millions of years. This fact was at odds with both the purely phyletic (figure 8.1B) and branching gradualist (figure 8.1C) models, which depict morphological change as a slow but consistent process through time. But then, what was the mechanism responsible for the remarkable stability associated with *Phacops* and so many other Paleozoic invertebrates?

This question points to an ambiguity in Eldredge's paper that would persist in the punctuated equilibrium literature for some time. Each of figures 8.1A–8.1C illustrates a particular model's implication for phyletic change. True to saltationist theories of evolution, figure 8.1A (the saltationist model) portrays evolutionary change as occurring only via evolutionarily instantaneous morphological leaps. When these leaps are not occurring, which is most of the time, the saltationist view tells us to expect no change. In accordance with phyletic gradualism, figures 8.1B and 8.1C both represent continuous morphological change in the phyletic mode, independently of speciation events.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

What about figure 8.1D, which Eldredge labels "Allopatric Model"? If we assume that this diagram is like the others in that it describes the particular model's commitments regarding the nature of evolutionary change along and between lineages, figure 8.1D seems to suggest that long-term stability is the expected consequence of the allopatric model. However, the theory of allopatric speciation itself makes no particular claims regarding what kind of phyletic change we ought to expect. In other words, Eldredge's figure 8.1D would have been just as faithful to the allopatric model if it had instead portrayed each lineage as undergoing steady morphological change through time, in addition to the "rapid bursts" that characterize allopatric speciation. The effect of this specific way of illustrating the "Allopatric Model" is to make it seem as though allopatric speciation can account for both of the significant features of trilobite evolution, which is not the case. In contrast, Eldredge himself asserts that the most probable explanation for the long-term stability is stabilizing selection—a mechanism that is entirely independent of allopatric speciation (Eldredge 1971, 160).

8.3.2 "The Famous Version of That Paper"

> The very next year the famous version of that paper, expanded with notions of species selection and other discussions, written by myself and Stephen Jay Gould, appeared in a book of collected original papers meant to stimulate paleontological research along more theoretical lines.
>
> —Description of Eldredge (1971) on Niles Eldredge's personal web page

Gould's first publication on punctuated equilibrium was "the famous version" of Eldredge's 1971 paper that he coauthored with Eldredge in 1972. The essay had been commissioned by University of Chicago paleontologist Thomas Schopf, who had invited Gould to participate in a symposium titled "Models in Paleobiology" at the 1971 meeting of Geological Society of America. As the invitation to participants makes clear, Schopf shared Gould's vision for a transformation of paleontology:

> What I would like the papers in this Symposium (and published volume) to accomplish is to identify and evaluate the theoretical models which are guiding (by accident or design) the development of various parts of our science. We now have both an extensive and a modern documentation of life in the past in the *Treatise* [*on Invertebrate Paleontology*], in discussions about specific groups in the *Journal* [*of Paleontology*], and elsewhere. The theoretical framework, however, dictates where one looks and how one goes about the descriptive process. . . . It

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> is also my conviction that the volume will strike a responsive cord [sic] in many of our colleagues who, having mastered the taxonomy of a group, now find it of interest to see additional ways in which their data can be significantly interpreted. In this way we can encourage the analytical "problem oriented" approach to paleontology. (Schopf to Raup, March 7, 1970, Schopf pap. 3, 30, Smithsonian Institution Archives, Washington, DC)

Gould's specific charge had been to present something on "models of speciation" (Schopf to Gould, March 9, 1970, Schopf pap. 5, 14). But, having little to say on the matter, Gould enlisted the collaborative effort of the younger Eldredge, who had recently submitted his manuscript to *Evolution* on precisely this topic.

The description of the Eldredge and Gould (1972) collaboration as "the famous version of" Eldredge's 1971 paper is accurate. The significant phenomena reported in the new paper were, again, the rapid morphological evolution and astonishingly long-term stability regularly observed in the fossil record. Again, the central argument of the paper is that the allopatric model of speciation is the best explanation for the observed morphological leaps. As in Eldredge's paper, the "unconscious guidance" of phyletic gradualism is singled out as the primary cause of paleontologists' failure to see how well these morphological breaks are explained if they are treated as real rather than as artifacts of an imperfect preservation process (Eldredge and Gould 1972, 108). To date, the principal scientific phenomena picked out by the phrase "punctuated equilibrium" have not been revised since Eldredge's paper.

In what ways do the two papers differ significantly? The most important difference where scientific theory is concerned is the explanation for long-term stability (renamed "stasis" in Eldredge and Gould [1972]). Whereas in the initial paper, Eldredge explicitly invokes stabilizing selection as the probable cause of stability, Eldredge and Gould repeatedly stress that both "*rapid* evolutionary events" and "a history of stasis . . . are among the expected consequences if most fossil species arose by allopatric speciation in small, peripherally isolated populations. This picture merely represents the application to the fossil record of the dominant theory of speciation in modern evolutionary thought" (108; emphasis in original). This revised view of allopatric speciation's consequences turns out to have major implications for the research program at a deeper, conceptual level. Whereas before Eldredge had invoked two mechanisms two explain two fossil phenomena—allopatric speciation to explain rapid evolution, stabilizing

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

selection to explain stability—we now have a single mechanism (allopatric speciation) to account for both. As mentioned above, in Eldredge's original paper, allopatric speciation might have been followed either by long-term stability or gradual phyletic evolution. It all depends on whether the speciation event is followed by a long-lasting stabilizing selection regime, a development that is neither here nor there from the perspective of allopatric speciation. By contrast, Eldredge and Gould (1972) argue instead that allopatric speciation can only be followed by "a history of stasis." By suggesting that both stasis and rapid evolution are necessary consequences of allopatric speciation, Eldredge and Gould effectively transform two distinct evolutionary phenomena into two stages of a single evolutionary process—the evolutionary process that produces punctuated equilibrium.

It would be noteworthy enough if this conceptual shift were all that followed from the revised view of the process of allopatric speciation. But there's more, and it is of considerably greater gravity. Remember Eldredge's reminder to his readers early in his 1971 paper that the vast majority of evolutionary biologists accept the allopatric model as the dominant, and plausibly the only, mechanism by which species originate. Combining this with the premise that allopatric speciation terminates in long-term stasis implies that evolutionary change only happens during speciation. This, in turn, yields the revolutionary proposal that all evolutionary change takes the form of punctuated equilibrium, and that *whatever explains punctuated equilibrium explains evolutionary change itself*. Now of course, everyone should be interested in a general explanation for evolutionary change. And no one interested in a general explanation of evolutionary change could possibly ignore punctuated equilibrium, because to understand punctuated equilibrium is to understand evolutionary change; after we've explained punctuated equilibrium, what's left? Thus, from a modest amendment to the theory of allopatric speciation comes an ostensibly knockdown argument for the importance of the paleontological perspective for understanding evolution. Only paleontologists are in a position to detect punctuated equilibrium, and punctuated equilibrium pretty much exhausts what there is to know about evolutionary change. As Gould and Eldredge would later confess, "we believe that a general theory of punctuational change is broadly, though by no means exclusively, valid throughout biology" (1977, 145).

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

8.4 Science Delayed (Again)

> For all the hubbub it engendered, the model of punctuated equilibria is scarcely a revolutionary proposal.
>
> —Stephen Jay Gould and Niles Eldredge, "Punctuated Equilibria" (1977, 117)

While Gould continues to defend the coherence and theoretical, biological, and disciplinary importance of punctuated equilibrium from beyond the grave (Gould 2007), the view itself evolved considerably from Eldredge and Gould (1972) and would never achieve the theoretical or disciplinary organizing role that he envisioned for it—even within the coterie of evolutionary biologists and paleontologists with whom it resonated. Our model of autonomous disciplinary emergence entails that the failure of punctuated equilibrium to provide a seat at the high table of evolutionary theory must be explained by at least one of two factors: either punctuated equilibrium (1) did not identify a set of related scientific problems whose solution achieved sufficiently high demand or (2) it was not something to which paleontologists had exclusive observational access. I submit that, being as it was a fossil pattern, paleontologists had a clear and arguably exclusive advantage with respect to authenticating the existence of punctuated equilibrium. So (2) cannot explain punctuated equilibrium's abortive disciplinary pretensions. How strong, then, is the case for (1)?

The case is very strong indeed. No matter which way we parse (1), it is not difficult to find a clear and convincing explanation for why punctuated equilibrium failed to serve as the foundation for a new evolutionary discipline. One way of not satisfying criterion (1) would be for punctuated equilibrium to actually pick out a set of related phenomena that were already well understood at the theoretical level. I would argue that most of the indifference associated with punctuated equilibrium is explained in this way. Whether we understand punctuated equilibrium as the deductive consequence of allopatric speciation (Eldredge and Gould 1972), or instead as allopatric speciation + stabilizing selection (Eldredge 1971), these are phenomena over which our theoretical grip is unshakably firm and had been for decades prior to Eldredge and Gould's work on the subject. One popular view, variously embraced by Eldredge and Gould themselves (Gould and Eldredge 1977, 117), is that both the phenomenon of impressive leaps across morphospace as well as its explanation in terms of the Wrightian

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

mechanism underlying allopatric speciation appear in all essential detail in G. G. Simpson's (1944) discussion of quantum evolution (mentioned in chapter 7; see also Simpson 1976, 5). Stabilizing selection had also been well understood for decades, being a straightforward consequence of Fisher's (1930) representation of natural selection as an optimization process that selects against deviations from optimum trait values.

Another way of not satisfying criterion (1) would be for punctuated equilibrium to not actually pick out a real phenomenon. There are at least two ways in which this interpretation can explain how punctuated equilibrium fell short of the high table. The first, already mentioned above, concerns the conception of allopatric speciation as necessarily entailing long-term stability. When both long-term stability and rapid evolution are depicted as consequences of allopatric speciation, it becomes possible to conceive of them as part of a single causal process. The phrase "punctuated equilibrium" picks out both stages of that causal process. But we do not have to accept the claim that evolutionary stasis is a consequence of allopatric speciation, for the simple reason that evolutionary stasis is not a consequence of allopatric speciation per se. Once we break the allegedly necessary connection between allopatric speciation and evolutionary stasis, we undermine the cogency of the claim that there is some causal process that has as its primary stages rapid evolutionary change and long-term stability. With the existence of that causal process in question, we're no longer compelled to see rapid evolutionary change and evolutionary stasis as causally related. And if they are not causally related, then it makes no more sense to say that there is something called "punctuated equilibrium" than it does to say that there is something called "brush driving" simply because I ordinarily drive my kids to school after brushing my teeth.

Another criticism of this kind raises the plausible worry that long-term stability in a lineage's fossil record is not sufficient to ground the claim that little to no evolutionary change is occurring within this lineage during that time (Futuyma 1987). In fact, pretty much everything we know about evolutionary change suggests that change was taking place in that lineage at that time and was simply not preserved in the fossil record—not all that surprising given that we all, paleontologists included, assume that there are entire lineages which were well represented while alive but none of whose members were lucky enough to fossilize. Thus, while we might agree with Gould and Eldredge (1977) that "stasis is data," morphological stability in

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

the fossil record is not sufficient evidence for stasis—at least, when interpreted as little to no evolutionary change.

8.4.1 Is Stasis Data?

> Stasis . . . is a phenomenon that calls for an explanation.
> —John Maynard Smith, "Paleontology at the High Table" (1984, 401)

Rejecting Eldredge and Gould's original interpretation of long-term stability in the fossil record will not make stasis go away. And after 40 years of sustained discussion in the evolutionary literature, it's not clear anything will. The long-term stability of interest in stasis suggests that many evolutionary biologists have deemed it to be a phenomenon worthy of considerable attention. Although the biological community rejects the thesis that all evolutionary change exhibits the pattern of rapid change followed by long-term stability (in other words, the thesis of Eldredge and Gould 1972), biologists from a range of disciplinary perspectives have attempted to explain this common and unintuitive feature of the fossil record. The most influential and widely accepted attempts have sought a biological explanation (rather than, say, a taphonomic one) for why the fossil record registers long-term morphological stability in so many species.

I would count this as a major achievement for Eldredge and Gould. To see why, recall our discussion of George Gaylord Simpson's complicated and clever relationship with the genetic evolutionary theory of the Modern Synthesis. Simpson argued that anyone who would take the population genetic theory of evolutionary change to be confirmed by fossil patterns implicitly affirmed the general principle that fossil patterns require biological explanations. In particular, the commitment to explaining fossil patterns biologically must be upheld even when those patterns violate the expectations of population genetics. It would be unscientific to use the fossil record entirely asymmetrically, able only to confirm what population geneticists already believe about the evolutionary process but never capable of suggesting the existence of additional evolutionary processes.

As compelling as this argument seems from the standpoint of scientific rationality, it was not until Eldredge and Gould's passionate campaign on behalf of punctuated equilibrium that a fossil phenomenon not predicted by population genetics gained wide acceptance as a biological phenomenon that required biological explanation. Now, as it turns out, the most

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

widely accepted explanation for stasis is essentially a population-genetic explanation, the argument being that we expect a novel trait to be eliminated from the species gene pool unless speciation occurs. Even though a species' range may be spatially and temporally structured in a way that exposes different local populations within that species to different selection regimes, allowing each local population to diverge from one another while they adapt to local conditions, the problem is that

> in the absence of speciation, much of the geographic variation we observe is ephemeral, leaving little imprint on evolution in the long term. This is chiefly because most local populations are ephemeral; over even moderately short spans of evolutionary time (tens or hundreds of thousands of years), the habitats to which populations are adapted shift, often over large distances, in consequence of climatic change. (Futuyma 1987, 467, quoted in Gould 2002, 802)

By contrast,

> reproductive isolation confers sufficient permanence on morphological changes for them to be discerned in the fossil record. Long-term anagenetic change [change along a lineage] in some characters is then the consequence of a succession of speciation events. (Futuyama 1987, 465, quoted in Gould 2002, 802)

Proponents of this explanation include Eldredge and Gould themselves (Gould 2002, 798–802; Eldredge et al. 2005).

According to this model, the fossil pattern called "stasis" is a by-product of the ephemerality of most of the evolutionary change not associated with speciation, rather than a specific stage of a distinct evolutionary process or the direct product of some causal mechanism. If that's true, Gould and Eldredge's (1977, 116) assertion "that each case of stasis has as much meaning for evolutionary theory as each example of change" now seems hard to accept. But not only is the model deflationary, it is also not a model whose design requires a deep understanding of paleontology—it follows straightforwardly from the statistical logic of population genetics. But that is not really the point. As the above quote from population geneticist John Maynard Smith illustrates, the major achievement for paleontology lay in the general recognition that this fossil phenomenon records a fact about the process of evolution. (Maynard Smith's comment is all the more significant in light of the fact that he cites stasis as one of two fossil phenomena to be credited with paleontology's graduation to the high table of evolutionary theory.) Of course, we can see now why we should expect stasis in the fossil record, given what we already accept about how evolutionary change

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

works. Thus, although punctuated equilibrium may not have been a conduit for disciplinary autonomy in the end, it should nevertheless be seen as a critical first step toward that agenda. After punctuated equilibrium, fossil patterns become genuine scientific problems whose solution requires evolutionary theory—new evolutionary theory, if it is needed.

Before moving on to describe Gould's third attempt to make a science, it's worth noting the similarities between Gould's advocacy of punctuated equilibrium and his plumping for the science of form. Recall that one of the distinguishing features of the science of form from a disciplinary perspective was its break with the inductivist and primarily descriptive practice that had characterized much of paleontology's history, a practice that Gould firmly believed had impeded scientific progress. The lack of progress, he argued, was due to the fact that descriptive generalization was not a fruitful lawmaking endeavor. The contingent nature of evolutionary history was such that attempting to generalize over species in the fossil record could not by itself disclose the evolutionary mechanisms responsible for producing fossil patterns. In contrast, the science of form prescribed an approach that designed causal models whose predictions could then be tested against the fossil record, a mechanism-disclosing strategy that Gould regarded as superior to inductivism. Both of these themes—anti-inductivism and the importance of modeling—would arise again in Eldredge and Gould (1972):

> No theory of evolutionary mechanisms can be generated directly from paleontological data. Instead, theories developed by students of the modern biota generate predictions about the course of evolution in time. With these predictions, the paleontologist can approach the fossil record and ask the following question: Are observed patterns of geographic and stratigraphic distribution, and apparent rates and directions of morphological change, consistent with the consequences of a particular theory of speciation? We can apply and test, but we cannot generate new mechanisms. (Eldredge and Gould 1972, 93)

Perhaps inspired by Schopf's comments in his invitation to Gould, Eldredge and Gould also took the opportunity to add a new dimension to the critique of inductivism that Gould developed in connection with the science of form. This was a broadside against the notion that the "inductivist philosophy that views 'objective' fact as the primary input to science," an "unsullied" descriptive approach to science that possessed the virtue of a theoretically neutral interpretation of data (Eldredge and Gould 1972, 84). Drawing on the current philosophy of science, they claimed that the view

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

of the inductivist as nature's doe-eyed registrar was a myth, and a positively pernicious one at that. By denying that the interpretation of data is always guided by some theory or other, the focus on allegedly dispassionate inductivism had blinded paleontologists to the fact that their own interpretations of the fossil record were themselves governed by particular theoretical commitments. When those commitments go unacknowledged, they tend to become more deeply rooted. Each new observation, guided as it is by theory, becomes yet another positive instance of that theory. Because these observation-guiding theories tell us what counts as data and what the data mean, the data themselves can never be brought into conflict with the theory; a feature of nature does not achieve the rank of "observation" unless it is (made to be) consistent with the theory governing observation:

> In paleontology, even the most "objective" undertaking, the "pure" description of fossils, is all the more affected by theory because that theory is unacknowledged . . . The inductivist view forces us into a vicious circle. A theory often compels us to see the world in its light and support. Yet, we think we see objectively and therefore interpret each new datum as an independent confirmation of our theory. Although our theory may be wrong. we cannot confute it. (Eldredge and Gould 1972, 86)

As with the unreliability of inductivism as a lawmaking strategy, the only refuge from the "vicious circle" of theory-laden observation was the conscious proliferation of models:

> To extract ourselves from this dilemma, we must bring in a more adequate theory: it will not arise from facts collected in the old way. Paleontology supported creationism in continuing comfort, yet the imposition of Darwinism forced a new, and surely more adequate, interpretation upon old facts. Science progresses more by the introduction of new world-views or "pictures" than by the steady accumulation of information. (Eldredge and Gould 1972, 86)

What Gould did not know at that time was that he would soon be party to introducing a world-view far more radical than punctuated equilibrium.

8.5 Third Attempt: Particle Paleontology

> How different then, is the real world from the stochastic system? . . . The answer would seem to be "not very."
>
> —Stephen Jay Gould et al., "The Shape of Evolution" (1977, 32)

The current resolution to the debate over stasis is symptomatic of how vulnerable paleontologists were to the threat of having their proprietary

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

phenomena explained (or explained away) by population genetic theories of evolution. Even in cases where paleontologists were able to garner widespread recognition of the evolutionary significance of certain kinds of fossil patterns, they were still forced to contend with the fact that the kind of explanation that describes replacement or maintenance of some alleles in local populations was the default form of solution to evolutionary problems.

There are a number of reasons for why population genetics was able to hold this default position. One is that population genetics had a history of success dating back the 1920s and was widely considered to be sufficient to understand the known evolutionary phenomena and, on that basis, evolutionary processes in general. A second reason is rooted in the fact that population genetics was effectively the only game in town. By the late 1940s there was no serious challenge to the theoretical picture provided by the Modern Synthesizers, and selection-based accounts of population change in particular had a solid enough mathematical and empirical foundation to command the status of a *normal form* of explanation of evolutionary phenomena—that is, the form that explanations of evolutionary phenomena were expected to take. Paleontology, interestingly, constitutes an exception to this generalization, where a willingness to appeal to alternative explanations, such as orthogenesis (directed evolution) and saltationism, persisted—perhaps not without reason—long after such appeals had become scientifically unacceptable in other biological disciplines (Simpson 1945; Brush 2009, 56–59). For those other disciplines, even if one were to find the Modern Synthesis picture to be an unsatisfactory model of some instance of evolutionary change, there simply was nowhere else to turn for an alternative general or well-understood form of solution to evolutionary problems.

The third reason for population genetics' default status—the explanation preferred by Gould—had to do with a general metaphysical view which holds that it is preferable for a phenomenon to be explained in terms of small things rather than sets composed of small things, because the small things are real and the sets they compose are not (or not as) real. Thus, from first principles alone we ought to accept that evolutionary change is best understood as changes in genes, the smallest biologically significant entities. In Gould's words, "reductionistic biases often lead us to seek an 'underlying' genetic basis for any overt phenomenon at any scale, and then to view data at this level as a fundamental locus for proper evolutionary explanations" (Gould 2002, 1219).

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

A successful campaign for the disciplinary autonomy of paleontology would need to confront each of the bases for the hegemony of population genetics, because each was by itself sufficient to justify that hegemony from a scientific perspective. Population genetics does resonate with reductionist intuitions in a way that might suggest to biologists that something like it must be basically correct, since it describes the mechanics of the lowest biologically significant level. Alternatively, given that it had been very successful at modeling the central evolutionary phenomena, at making predictions, and at solving outstanding biological puzzles, biologists would have been justified in pursuing it to the exclusion of rival models. Were it to have had neither of these highly salutary qualities, however, there would still have been a clear rationale for its theoretical dominance, because there was no alternative, let alone one of equal precision and generality.

This, then, would have to be the starting point for paleontology as an autonomous theoretical discipline: develop a new form of solution to evolutionary problems. Only then would it be possible for paleontologists to independently satisfy the demand generated by their disclosure of unexpected evolutionary phenomena recorded in the fossil record. Without a means of independently satisfying that demand, paleontologists would be forced to cede the real scientific work of solving evolutionary problems to those who understood how to do it.

Yet even with an alternative form of solution in hand, our model of disciplinary emergence entails that paleontologists still would need to identify a set of evolutionary problems that their newly developed solution form could potentially resolve. In this section I argue that Stephen Jay Gould played a unique and pivotal role in the philosophical and biological justification of a form of evolutionary problem-solving that stands as a genuine alternative to population-genetic explanation and was similarly pivotal in articulating the sorts of biological phenomena to which paleontologists might apply this problem-solving form. But although this distinctive approach to explaining evolutionary phenomena was firmly in place by the early 1970s, it would take more than a decade for it to be recognized as such outside the paleontological community. As predicted by our model, this was due to the fact that there were not yet any evolutionary problems raised by the fossil record whose solution was in sufficiently high demand. The new form of solution—"species as particles"—was a solution without a problem.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

8.5.1 Paleontology's Task

> We are convinced that sequences of unique historical events have strong general components (regulated by laws that are independent of time, space and taxonomic group)—and that it is the (heretofore neglected) task of paleontology to discover them (not by induction from empirical catalogues, but by attempts to model results with relatively simple systems).
>
> —David M. Raup and Stephen Jay Gould, "Stochastic Simulation and Evolution of Morphology" (1974, 321–322)

Let us first recall Gould's argument from Dollo's Law to the imperative of a model-centered paleontology, given the goal of finding evolutionary laws. Dollo's Law—that evolution is irreversible—was based on the notion that particular phenotypes are the result of a long chain of highly improbable evolutionary events. Trying to discover evolutionary laws by examining the phenotypic features of the fossils at the ends of different Markov chains is a doomed strategy, because each of these features would have been otherwise were we to replay the tape of life, and evolutionary laws must express facts that *could not have been otherwise*—"a kind of timeless generality" that would be the same no matter how many times we replayed life's tape. The only fruitful lawmaking strategy would be to construct models of evolutionary processes and then fit them to fossil patterns: laws "must be based on immanent processes that produce events, not on the events themselves" (Gould 1970a, 209).

In 1970, Gould clearly had a sense that the nature of history itself was the major obstacle standing in the way of paleontology's emergence as a theoretical discipline. Individual phenotypes record so many chance events, so much evolutionary noise, that paying particularly close attention to them blinds us to the "nomothetic undertones"—the "timeless generalities" that govern evolution. Both his work on allometry and on punctuated equilibrium can be understood as attempts to get around history's vicissitudes by looking for patterns that should obtain independently of any particular historical factors. No matter what happens during the history of a lineage, each species, no matter what it is like, will be forced by natural selection to buckle under the selective pressure of the power function. No matter what happens during the history of a lineage, each species, no matter what it is like, will only experience evolutionary change by rapid leaps across morphospace followed by millions of years of stasis. The message underlying both research programs is that there are evolutionary processes whose

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

trajectories are not sensitive to the effects of history on phenotypes. From the perspective of these evolutionary processes, all species are equal.

Around the time that Gould was articulating this view about evolution, Thomas Schopf was carving out a different path to the same conclusion. In section 8.3.2 I mentioned Schopf in connection with his having organized the symposium on Models in Paleobiology, which produced Gould's first paper on punctuated equilibrium. That paper appeared in 1972 in a volume edited by Schopf, who wrote an introductory essay spelling out the form of and motivation for aggressively pursuing a model-based approach in paleontology. Schopf was deeply impressed with the immigration-extinction equilibrium approach to island biogeography recently developed by MacArthur and Wilson (1963, 1967; discussed in Sepkoski 2012, 222), and in his essay advocated this approach as an alternative or complement to the traditional, inductivist activities of "description, illustration, and cataloging of the diversity of ancient life" (Schopf 1972, 9).

Although he shared with Gould the desire to redistribute research efforts in paleontology to increase theoretical work, his reasons were rather different, at least measured against those articulated in Gould's "Dollo's Law" paper. Schopf had been concerned, both in his introduction to the "Models" volume and elsewhere ("Talk to Faculty, November 10/1972," Schopf pap. 2.21), with the question of "whether paleontology is obsolete," or whether instead there was anything that paleontology could contribute to our theoretical understanding of the evolutionary process strictly by virtue of its focus on large time scales. Now, this was not a novel question. Simpson, for example, had asked and answered it in the affirmative in the 1940s, and it had occupied plenty of paleontologists in the interim. But Schopf's reasons for affirming the importance of the paleontological perspective do stand out as particularly novel. What he suggests in his "Models" introduction is that the effect of shifting time scales is to make different units biologically significant.

Figure 8.2 is Schopf's (1972) diagram representing different biological units in time and space. Here he includes two alternative representations of the concept of a *community*, meant to illustrate the way in which different time scales render different units biologically significant:

> If one views communities from a historical point of view . . . then *the particular species composition of a given community is emphasized*. In that case one may refer to several types of level-bottom communities according to the dominant members;

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

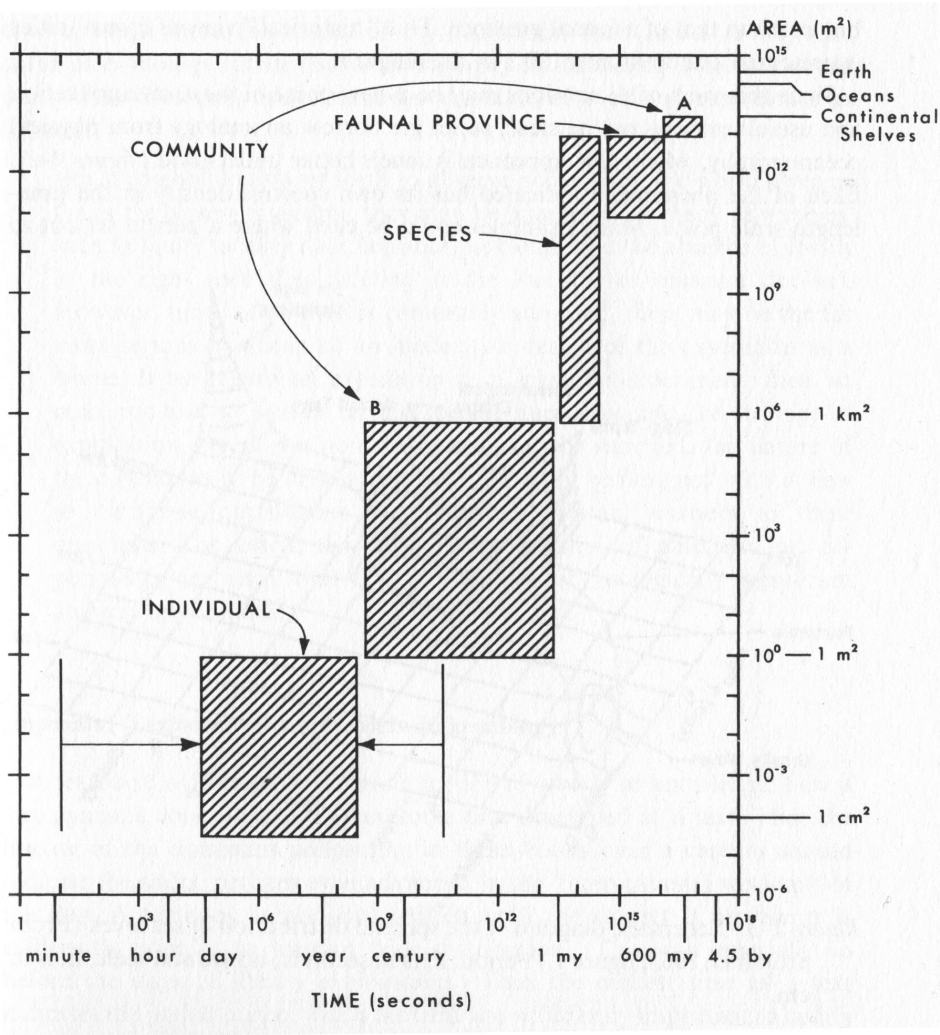

Figure 8.2

Schopf's diagram from the 1972 *Models in Paleobiology* volume. His caption reads, "Dimensions of paleontologically important items with time in seconds and area in m^2. Community 'A' is based on an equilibrium model of communities; Community 'B' is based on an historical model of communities."

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> these may change from place to place. *Such communities would exist on the time scale of the life of that assemblage*. This time is greater than that of an individual but less than the time it takes to evolve a new species because, strictly speaking, the substitution of new species changes the community. (Schopf 1972, 18–19)

The key idea here is that, on the time scale associated with a species-centered conception of community (i.e., B in figure 8.2), the histories of lineages are deeply significant, because individual species record sequences of evolutionary events, and there is nothing to the community over and above the assemblage of those particular species.

However, the locus of biological significance shifts if we move to the time scale associated with a conception of community that emphasizes "particular 'general environments'":

> In one sense, communities such as the level-bottom community, the interstitial community, or the reef community transcend faunal provinces even though the particular species composition may change. In this view, communities have existed as long as particular "general environments" have existed (from 0.1 to 1 billion years) over areas as large as most of a continental region. (Schopf 1972, 18)

Schopf described this as the "'equilibrium' view" of biological communities, where the biologically significant entity is the "general environment," which persists across geologically long intervals of time and whose identity is decoupled from the particular species that inhabit that environment. Schopf's favorite analogy for the equilibrium view was that of "a gas law in which the state of any particular molecule is immaterial to the general description of the behavior of the volume as a whole." In contrast to the historical view of communities, here "the particular history of species may be immaterial . . . in some important senses, every species is 'equally good'" (12–13).

The critical complement to these *principled* arguments for the irrelevance of species' history was provided by evolutionary polymath Leigh Van Valen, whose study of extinction rates appeared to show that, in the case of extinction, a species' history actually is irrelevant. Self-published in 1973, Van Valen's "A New Evolutionary Law" overturned the conventional view that a taxon's probability of extinction increases with its age by showing that "all groups for which data exist go extinct at a rate that is constant for a given group"—a phenomenon he dubbed "The Law of Extinction."

Here is how his argument works. Take any taxonomic group (say, a family). Calculate the longevity for each subgroup by measuring the distance

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

in years from its origination time to its extinction time. Now plot each of these subgroup longevities starting at year 0, so that all subgroups exist at year 0 and no subgroups exist after the extinction of the longest lived subgroup. The conventional expectation was that, as we get further and further away from year 0, we should see an increase in the percentage of subgroups that go extinct in a given time interval. This corresponds to the view that a taxon's probability of extinction increases with age: when taxa are young, their extinction probability is low, and so years nearer to 0 should see relatively small percentages of subgroups lost to extinction. As taxa get older, their extinction probability rises, and so subsequent years should see subgroups lost in relatively greater proportions. The extinction probability can thus be thought of as the "mortality rate" for a given age. What Van Valen found was that the percentage of extinctions per unit age remained constant across ages within a group—for all groups examined (figure 8.3). If the frequency distribution of extinction events is the same all ages, then each age has the same mortality rate. Thus, the probability of extinction does not vary with age. With respect to taxon age, then, extinction is "a randomly acting process" (Van Valen 1973, 17). (An analogous situation

TABLE 1. Life table for hypothetical case of linear survivorship.

Age at start of interval x	Number dying in age interval d_x	Survivors at start of interval l_x	Mortality rate q_x
0	4000	10000	0.4
1	2400	6000	0.4
2	1440	3600	0.4
3	860	2160	0.4
4	526	1300	0.4
5	312	774	0.4
6	186	462	0.4
7	114	276	0.4
8	66	162	0.4
9	42	96	0.4
$\geqslant 10$	54	54	—

Figure 8.3
Actuary table from Raup (1975).

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

would be a case where we found that the human mortality rate was the same for all ages. If the probability of dying is the same no matter how old one is, then death does not discriminate between ages [although it might still discriminate based on other factors, like diet, for example].)

Van Valen's "astonishment" in the face of the Law of Extinction compelled him to reflect on its significance for the investigation and interpretation of evolutionary processes in general. In particular, he observed that his extinction results suggested a fundamental difference between the evolutionary processes operative at this level versus those operating at the genetic level:

> We see here a major difference from the usual theory of genic selection. The latter depends only on the current distribution of frequencies of alleles and their interactions with each other and the environment. It does not depend at all on the process by which the current distribution was obtained. In formal language, it is a Markov process. But any process can be made Markovian by choice of a suitable level of analysis. With extinction we can see that a non-Markovian analysis of selection is appropriate. This is probably true for the general case also. (Van Valen 1973, 18–19)

When viewed from the genetic level, our understanding of a species' evolutionary trajectory depends solely on its genetic constitution and that constitution's relationship to the environment. That genetic constitution is a record of the species' evolutionary history, a history that, as Gould observed, is a highly contingent matter of fact. By contrast, Van Valen's "Law of Extinction" suggests that where a species is going does not depend at all on where it has been in certain respects. Because extinction is a "randomly acting process," nothing in the evolutionary history of a species can prepare it for what's to come.

8.6 The Marine Biological Laboratory

> You are right in arguing that we are among the first group of scientists to use a consciously nomothetic approach to paleontology (there's that word again, but it's the right one and I don't apologize). . . . The more you think about it, the more you realize how heretical the search for a kind of timeless generality appears within a science so deeply committed to historicity. . . . The only reason we may not cause the stir we should is that paleontologists are so notoriously unaware of the philosophical implications of their methodologies.
>
> —Stephen Jay Gould to Thomas J. M. Schopf, October 5, 1973, quoted in D. Sepkoski (2012, 236)

Gould had argued in 1970 that species differences disguise lawful evolutionary processes, and so should be ignored when searching for evolutionary

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

laws. Schopf had argued in 1972 that species differences should be ignored when working with time scales that render them biologically insignificant. In suggesting "that extinction within a group behaves like radioactive decay" (Raup 1975, 82–83), Van Valen's empirical study (which had "been circulating in *samizdat* since December, 1972" [Van Valen 1973, 21]), appeared to show that, in the case of extinction, traditionally important differences between subtaxa—including species—should be treated as irrelevant because they actually are irrelevant. It was against this conceptual backdrop that Gould met with David Raup, Daniel Simberloff (a newly minted E. O. Wilson PhD student), and Schopf, who organized their gathering in August 1972 at the Marine Biological Laboratory (MBL) in Woods Hole, Massachusetts.

In the remainder of this section I will argue that the MBL group was solely responsible for transforming the philosophical and biological insights of Gould, Schopf, and Van Valen into a general form of evolutionary explanation, and that it is only after MBL that a genuine alternative form of solution to evolutionary problems emerges as part of the standard evolutionary toolkit.

8.6.1 "A Comparison of Real and Random Clades"

> Paleontology has traditionally been focused on the idiographic (why did *this* crinoid become extinct at *that time*); many of its unsolved dilemmas and some of its current sterility may reflect this powerful bias of methodology rather than any limitation of available information. We prefer the nomothetic approach. We shall try to abstract the common elements hidden beneath the bewildering and all but impenetrable verbiage of our nomenclature and stratigraphy by constructing a model that makes no reference to the peculiarities of individual taxa.
>
> —David M. Raup et al., "Stochastic Models of Phylogeny and the Evolution of Diversity" (1973, 526)

This passage, taken from the first pages of the first paper published by the MBL group, encapsulates the deeper motivation and strategy guiding the MBL work. It also reflects the particular influence of several of the views expressed by Gould in earlier work. Recall our initial discussion of his first published essay, "Is Uniformitarianism Necessary?" (1965), in which he had argued that to be a science was to produce laws that one could project onto unexamined cases. Here again in the explanation for the MBL project we see an explicit connection drawn between scientific progress within paleontology on the one hand and need to formulate laws of nature on the other. Additionally, the passage asserts that "the nomothetic approach"

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

requires that we "abstract the common elements . . . by constructing a model that makes no reference to the peculiarities of individual taxa"—precisely what Gould had argued in his 1970 "Dollo's Law" paper. If the highly idiosyncratic phenotypic outcomes of the evolutionary process are impeding our ability to discover evolutionary laws, we need to find some way of disclosing laws that does not rely on phenotypic outcomes.

The idea of modeling evolutionary processes without making "reference to the peculiarities of individual taxa" was at the core of MacArthur-Wilson work on island biogeography that had so influenced Schopf. MacArthur and Wilson wanted to see what you could predict about species diversity (measured in number of species) on an island if you assumed nothing about individual taxa beyond the fact that they arrive on an island at some point and go extinct at some later point. Given these extremely conservative assumptions, they proposed that species diversity on an island can be modeled as an equilibrium process governed by the relationship between the rate at which species are migrating to the island and the rate at which species are going extinct on the island. Conceived of in this way, the number of species on an island is the point (*s*) at which the migration rate equals the extinction rate—the "saturation" point. At this point, species diversity is at equilibrium: if the island's rate of extinction exceeds the rate of migration, the number of species on the island drops, which in turn causes migration to accelerate. If it continues to accelerate beyond the rate of extinction, the island will be oversaturated, at which point, the extinction rate begins to increase, bringing the number of species back to the saturation point—the equilibrium point where the extinction rate equals the migration rate. "Therefore," they argued, "in order to predict the number of species on an island we need only construct these two curves and see where they intersect" (MacArthur and Wilson 1963, 377; see figure 8.4).

This was an abstraction from peculiarities of individual taxa par excellence because, according to the model, we can know to a good approximation how many species are on an island simply by knowing the rate at which species are coming and going, so to speak. We need know nothing about their evolutionary history, their phenotypic properties, or their interactions with one another. On this model, species are represented simply as biologically undifferentiated particles bouncing around in ecological space, determined by general equilibrium forces to settle at a particular level of diversity. Impressed by this approach, the MBL group sought a "continuation of the conscious application of equilibrium models to paleontological data, asking

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

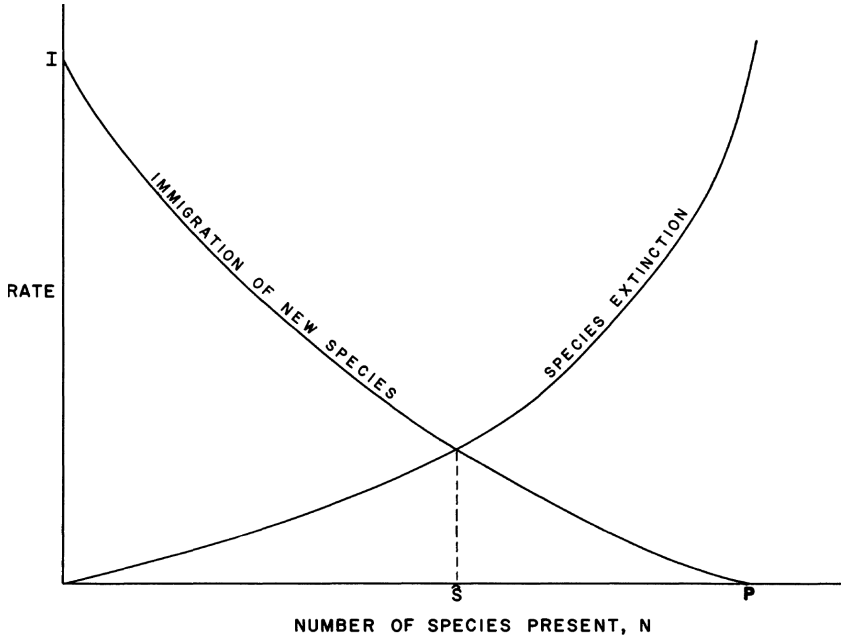

Figure 8.4
An equilibrium model of the interaction between immigration rates and extinction rates, from MacArthur and Wilson (1963).

simply, 'Can the processes resulting in local equilibria in ecologic time also be used to predict events in evolutionary time?'" (Raup et al. 1973, 525).

In paleontological terms, this meant asking first "whether an equilibrium pattern of phyletic radiation can exist?" (Raup et al. 1973, 525). To answer this question, the MBL group used a computer to simulate what pattern diversity might show across evolutionary time if lineages had no biologically significant properties aside from an origination point and an extinction point. In basic form, their program simulated an evolutionary process in which, at every time step, each lineage had an equal probability of doing one of three things: (1) go extinct, (2) produce another lineage and persist unchanged to the next time step, or (3) persist unchanged to the next time step without producing another lineage. As the group stressed, these

> constraints make no assumptions about the evolutionary mechanisms involved, such as natural selection and mode of speciation. They only specify the branching format by which evolution proceeds. Moreover, genetic or morphological change plays no explicit role in the formulation. (528)

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

From this perspective we see evolutionary lineages as MacArthur and Wilson envisioned species on islands—formless, physically homogeneous particles aimlessly meandering across evolutionary time, individually subject to random origination and extinction. Under these conditions, where species are mere particles, what sort of diversity patterns would we expect to see?

The first item that caught the MBL group's attention was the fact that the simulated random clades tended to exhibit an enormous range of sizes and shapes. Conceived of as mere particles, lineages could be expected to take on any diversity pattern you like (with a couple of notable exceptions—see below). As the authors pointed out, this result was of both biological and methodological significance:

> When faced with such variation in evolutionary patterns, paleontologists are inclined to suspect or even to postulate that the organisms involved are inherently different—that the various taxonomic groups differ from one another in evolutionary potential because they differ in population structure, reproductive systems, mutation rates, dispersal systems, and so on. But the simulation modeling shows that two or more groups operating under identical constraints—that is, having the same evolutionary potential—can behave very differently. We do not rule out that such differences in the real world may be due to inherent biological differences. We only contend that the observed variation does not in itself demand such an explanation. Above all, the simulation should be a warning against using patterns of diversity as the major evidence for differences in evolutionary potential. (Raup et al. 1973, 534)

The principal point here is that if we can generate differences in clade diversity patterns in the absence of differences between their constituent lineages, the differences in clade diversity patterns do not necessarily require explanations based on "the peculiarities of individual taxa."

More important than its ability to construct a variety of diversity patterns was the random model's ability to generate patterns that were astonishingly similar to those exhibited by real clades. In real clades, these patterns had traditionally been understood to result from the long-term directional influence of natural selection. But over the course of four eye-opening articles, the random appearance and disappearance of lineages depicted by the MBL simulations would "display most of the ordered features generally associated with uni-directional selection" (Raup and Gould 1974, 305).

For the MBL group, there were essentially two take-home messages with which the simulations could be credited. The first was a methodological point about traditional scientific practice within paleontology on the one

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

hand and what kinds of evidence warrant inferences regarding the influence of natural selection on the other:

> Modern paleontology has retained a vestige of idealistic morphology in its traditional argument for the role of directional causes in macroevolution. The presence of order in the results of evolution has been taken uncritically as definite evidence for the production of such order by directed causes. The results of evolution are the phylogenetic trees that paleontologists love to draw. The marks of order include evolutionary trends, correlation between character, and morphological coherence of taxonomic groups. The postulated cause, for the last thirty years at least, has invariably been uni-directional selection. We have come to doubt that the formal pattern of change is an adequate argument for directed causes. (Raup and Gould 1974, 305)

What the simulations imply for the working paleontologist was, in the succinct language of Raup and Gould, that "an orderly pattern . . . supplies no proof for uni-directional selection if that pattern can be generated by random processes as well."

The second, more substantive implication was that "life's history . . . might well be ['ontologically random'] in several important respects" (Gould et al. 1977, 23). Unsurprisingly, the MBL group saw the fact that real and randomly generated clades resemble one another in many important respects as partly supportive of the view that real clades are randomly generated. In other words, the evolutionary processes that produce real clades are processes for which "the peculiarities of individual taxa" do not matter. And if the peculiarities do not matter, then those evolutionary processes really do see species as particles—a view Raup would affirm nearly 40 years after the formation of the MBL group: "And so, I guess, I was arguing then and I [hold] still that at certain scales, particularly, over large [chunks] of time that it comes out indistinguishable from random. Therefore, species are particles" (David Sepkoski, interview with David M. Raup, July 20, 2005).

8.7 Science Delayed (Yet Again)

> Recently . . . a far more radical suggestion for a nomothetic paleontology has been advanced. The traditional view admits intrinsic historicity of events and applies nomothetic principles to abstracted aspects. The radical view asks if there might not be a sense—at an appropriate scale of analysis—in which the events themselves are essentially nomothetic in character. That is, might a biological object be treated without reference either to its taxon or to the time in which

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> it lived. Might species be like the molecules of classic gas laws—might they be treated as particles fully comparable one with the other no matter when they actually lived, or what their size, shape, and phylogenetic affiliation? Might a Tertiary clam and an Ordovician nautiloid be reasonably treated as two molecules of hydrogen diffusing from a vessel?
>
> —Stephen Jay Gould, "The Promise of Paleobiology as a Nomothetic, Evolutionary Discipline" (1980, 113)

The MBL work produced a number of fascinating results and attracted an appreciable amount of scientific attention. But the joyride was about to end. Coming down the pike was a critique initiated by paleontologist Steven Stanley that would neuter an entire class of MBL claims regarding similarities between real and random clades. Published shortly after the Gould paper quoted above, Stanley's critique argued that MBL's impressive variety of diversity patterns was not attributable to the random nature of originations and extinctions, but rather to the fact that MBL clades had "an average standing diversity of about 4 'lineages'"—unnaturally small from an empirical standpoint (Stanley et al. 1981, 17). Generally speaking, small populations are much more sensitive to the effects of chance than are larger populations. In the context of the MBL random model, this means that the simulated clades will tend to exhibit far greater instability in their number of lineages, expanding and contracting much more wildly than would a population of, say, 25 lineages. (Here's an example: suppose that each lineage in a clade has a 50% chance of branching at any given time step. For a clade with only four lineages, the probability of that clade doubling in size in one time step is $.5^4$, or about .06. For a clade with 25 lineages—which is closer to observed sizes—the doubling probability is $.5^{25}$, or .00000000298. The probabilities become even more divergent as the chances of branching drop). Therefore, Stanley argued, the radical instability of the MBL clades should be attributed to the size of MBL clades, not their random behavior.

Now, because there is no reason to think that real clades are in general this small, there is no reason to think that the enormous variety observed in real clades is due to the random nature of origination and extinction. Stanley et al. (1981) confirmed that this was the case by running the MBL simulations with the same chances of branching and extinction but with more realistic clade sizes. Their simulations generated diversity patterns that were consistently more stable and more uniform. With reasonably sized random clades exhibiting a pattern of stability that looked very unlike the

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

distribution of patterns in nature, Stanley et al. argued that the influence of natural selection was still the best explanation for differences in diversity patterns among real clades. Thus, contrary to the stern warnings made by the MBL group—particularly Gould et al. (1977)—the sorts of diversity patterns traditionally associated with natural selection could not have been generated by random processes, and thus retained their standing as evidence for the causal influence of natural selection. Obviously, this critique also calls into question the plausibility of the "species-as-particles" thesis. If

1. differences in clade diversity patterns are best interpreted as differences among the constituent lineages with respect to their chances of branching and extinction, and
2. differences in these chances depend on differences in the "peculiarities of individual taxa," and
3. the peculiarities of individual taxa affect the diversity patterns of the clades in which they are situated, then
4. there is some evolutionary process involved in the generation of clade diversity patterns for which species are not particles.

The failure of MBL to generate disciplinary autonomy is easy to understand if we adopt the more conservative, "null model" perspective for which the group often argued. If the MBL work was not intended to identify or solve any real scientific problems, but instead just to make an adjustment in the way that paleontologists made inferences from fossils, then there would obviously have been no basis on which disciplinary autonomy could have rested, and thus no special problem of MBL's failure to produce it. But I personally do not buy the claim that the chief intent of the MBL group was to raise the evidentiary burden for supporting claims about the influence of natural selection. The MBL group was overtly committed to the formulation of explanatory laws of nature that operated "independent of time, space, and taxonomic group." Each of the MBL articles puts this lawmaking strategy on prominent display. This makes an ontologically rich interpretation of the MBL mission much more credible.

I have included the MBL episode in this discussion partly because of Gould's central role in various dimensions of the project, but also because I believe that, had it possessed all of the features that our philosophical model requires for the emergence of an autonomous discipline, the MBL work would have been sufficient to act as midwife for the birth of a new

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

discipline. It was sufficiently innovative, sufficiently clear, and sufficiently general and important to have generated the requisite demand on the part of the biological community. But the historical record indicates that this did not occur, a fact that can be understood from the perspective of our model.

Drawing on the dynamics of branching described in chapter 5, I argued in chapter 7 that disciplinary autonomy requires (1) the identification of a set of related problems whose solution is in sufficiently high demand and (2) the availability of a privileged perspective, uniquely well suited for solving that set of problems. The phenomena with which the MBL group were engaged could only have been appreciated through an examination of the fossil record, and so the MBL work has satisfied criterion (2). This means that either the MBL group did not identify a set of related problems or that solutions to the problems they were able to identify were not in sufficiently high demand (or both).

Let's start by dividing the MBL results into two principal categories: clade diversity patterns and morphological trends. One of the great innovations of the MBL group was to treat the extreme variability of clade diversity patterns as a specific evolutionary phenomenon, one that might be predicted from certain assumptions about what sorts of things can count as biologically significant entities on which evolutionary mechanisms could operate. More specifically, the MBL group in many places suggests that there might be evolutionary processes for which species differences do not matter and that variability of clade diversity patterns is a direct consequence of those processes. By showing that a suitably adjusted MBL model did not predict such variability, Stanley et al. (1981) illustrated that the MBL results did not support the claim that the variability of diversity patterns was a real evolutionary phenomenon per se rather than a by-product of different adaptive constraints operating at different times and places on the peculiarities of individual taxa—the traditional view. In this way, the fate of the MBL claims about variability are much like Eldredge and Gould's claims about punctuated equilibrium. Critics of punctuated equilibrium argued that stasis is not a real phenomenon because there is no evolutionary process that has stasis as a direct consequence. Instead, stasis is simply a by-product of the interaction of a few population genetic processes. Similarly, the MBL group suggested that variability of clade diversity patterns is a real evolutionary phenomenon because there is some specific evolutionary process—random origination and extinction of lineages—that predicts that variability even

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

in the absence of species differences. Stanley et al. undermined the argument for the existence of this process, thereby undermining the claim that variability is a real phenomenon. In this respect, the MBL group failed to identify a real scientific problem.

What about the MBL results concerning morphological trends? Unlike variability in diversity patterns, the biological reality of morphological trends was not in question. But then, neither was it poorly understood. The existence of morphological trends is straightforwardly predicted by the theory of natural selection, which describes the gradual modification of organisms in a population as they ascend an adaptive peak—a peak fixed by the selection pressures operative in that population's environment. This is a process for which we have a detailed mathematical understanding at the genetic level, with different alleles being associated with different morphological traits that affect probabilities of reproductive success in a given environment. Thus, a general solution to the real problem of morphological trends was not in high demand.

8.7.1 A New Form of Solution to Evolutionary Problems

Despite its failure as a purported revolutionary discovery, punctuated equilibrium still holds the title honor of being a fossil pattern that was not predicted by population genetic theory but was recognized as a biologically significant phenomenon that required an evolutionary explanation. In that, it represents a major achievement for paleontology. In the remainder of this section I will argue that what was significant about the MBL project was that, despite its failures, it firmly established the availability of an alternative form of solution to evolutionary problems.

The chief mechanism by which the MBL group was able to achieve this result was through their simulations' ability to display the raw explanatory power of an entirely new style of evolutionary problem solving. Leading up to Stanley's critique, the MBL group could legitimately claim that treating species as undifferentiated biological blobs allowed one to simulate "most of the ordered features generally associated with uni-directional selection"—in other words, most of the evolutionarily significant features of the fossil record. It is probable, as Huss (2009) has argued, that the specific graphical form in which the MBL results were presented had some positive effect on the degree to which paleontologists were disposed to accept the MBL results. However, in my estimation the most impressive

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

and convincing aspect of their study was the MBL model's ability to accommodate the majority of evolutionarily significant fossil phenomena. Now, of course, the claims about clade diversity pattern did not survive scrutiny, and the claims about the causes of morphological trends were not a strong competitor to the traditional population genetic explanation. But in a way it did not matter, because the MBL simulations showed that species-as-particles was the kind of account that, if true, could explain the phenomena at the heart of evolutionary paleontology. These were the phenomena that formed the evidentiary basis for paleontology's participation in the Modern Synthesis. Nothing could be more significant to the potential success of an alternative form of solution than a clear demonstration of the explanatory utility of that approach. Table 8.1 summarizes the major results of the MBL collaboration.

This brings us to the second strength of the species-as-particles approach to solving evolutionary problems. The MBL project could not have helped but turn paleontologists' attention toward the possible future of their discipline. The MBL papers themselves did not have to produce valid scientific results in order to alert the paleontological community to the potential fruitfulness of both (1) the conceptual shift away from the significance of species differences and toward that of higher taxa and (2) the consequent ability to make predictions about what we ought to see in the fossil record. This would have been particularly important for the goal of developing an alternative way of solving evolutionary problems—they were making quantitative predictions about the properties of clades, for crying out loud!

Most critically, the species-as-particles approach allowed MBL to make predictions about "the shape of evolution" that could not have been made with population genetic theory. Regardless of whether these predictions turned out to be true, the mere ability to make them would have established the species-as-particles approach as both a source of new evolutionary problems—a new strategy for suggesting and identifying genuine evolutionary phenomena that population genetics was powerless to explain—and a unique form of solutions to those problems, one that could legitimately claim explanatory hegemony over a particular domain of biological facts. Under the kind of paleontology suggested by the MBL project, the solutions to those evolutionary problems—old as well as new—would take the form of laws of nature that operate "independent of time, space, and taxonomic group." Because of their freedom from the "peculiarities of individual taxa,"

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Table 8.1
Summary of Similarities and Differences between Real Clades and MBL Simulated Random Clades (year of publication in parentheses)

Real clades	Simulated random clades	
X	X	Spindle-shaped (1973)
X	X	Origins concentrated in early part of sequence (1973)
X	X	Period of nonextinction early in time span (1973)
X	X	Period of nonorigination late in time span (1973)
X	X	Highly correlated evolution of characters (1974)
X	X	Long-term morphological trends (1974)
X	X	Morphological coherence of Monophyletic Groups (1974)
X	X	Incomplete filling of morphospace ("clumpiness") (1974)
X	X	Regular "unfolding" of morphology (1974)
X	X	Specialization of derived forms (1974)
X	X	Evolutionary rates vary directly with morphological complexity (1975)
X	X	Alive clades larger than extinct clades (1977)
X	X	Extinct clades are most diverse in the first half of their lives (1977)
X		Morphological convergence (1974)
X		Low values of uniformity for clades of wide taxonomic separation (1977)
X		Mass extinction (1973)
X		Very rapid radiation (1973)
X		Persistence of single-lineage clades (1973)
X		Highly variable clade size (1973)

not to mention of time and space, these laws would describe evolutionary processes that were thoroughly insensitive to the contingency inherent in the phenotypic level.

The elegance and simplicity of the MBL models would have been further evidence of the potential fruitfulness of the MBL approach. Recall that Gould had trumpeted the fecundity of modeling not only because of its ability to disclose the existence of laws of nature, but also because sufficiently simple

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

models were more conducive to understanding and explanation, and they often had a prodigious range of applicability. In the case of the power function, Gould advocated its use not because it was the most accurate quantitative description of form (because it was indeed not the most accurate) but because it was simple, general, and deeply intuitive on a biological level. Because of its simplicity, its relatively few and relatively well-understood parameters, and its freedom from the phenotypic detail of specific taxa, the MBL model was well positioned to be extended to the full range of fossil trends, as well as to suggest new kinds of fossil trends that we could plumb the stratigraphic depths in search of. Gould had finally found the fruitful lawmaking strategy he had been seeking. In the more general context of this book's argument, we can say this: Gould's metastrategy for producing a distinctively paleontological Darwinian population had begun to pay off.

With an alternative form of problem-solving firmly in hand, the task that lay ahead was to identify a set of fossil phenomena that could not be predicted or explained by population genetics but was of unimpeachable evolutionary importance. That set of phenomena would be the corpus of regularities associated with mass extinction.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

9 Wonderful Death

> This brings me to what I see as the greatest impact that palaeontology is having on the way we see the mechanisms of evolution. We have been familiar for a long time with the dramatic disappearance of the Dinosaurs at the end of the Cretaceous. It is now apparent that massive extinctions, involving many different taxa, have been a repeated feature of evolution.
>
> —John Maynard Smith, "Paleontology at the High Table" (1984, 402)

> But the excitement of paleobiology in the world of ideas lies in the theories it has developed at the second tier of sustained trends during millions of years in normal times and the third tier of mass extinction. We have reclaimed our proper role among the evolutionary sciences—as generators of testable and irreducible theories about pattern in the history of life, not passive recordkeepers of a phenomenology fully explained at the first tier.
>
> —Stephen Jay Gould, "The Paradox of the First Tier" (1985, 10)

By 1985, Stephen Jay Gould could triumphantly and correctly claim that paleontologists had "affirmed the theoretical independence of [their] discipline," a view that he shared with population geneticist John Maynard Smith (Gould 1985, 11). In addition, both men recognized the pivotal role of extinction studies in particular in making this claim credible. How did this happen? According to chapter 5's discussion of how new disciplinary Darwinian populations emerge, three conditions would need to have been satisfied. First, extinction would have to be a real phenomenon of general evolutionary significance, into which paleontologists had unique insight; this helps to satisfy the *intrinsicality* condition. Second, conceived of as scientific practices, both the theoretical problems associated with extinction and the unique tools developed to solve them would need to be passed

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

down from "parents" to "offspring"—that is, "apprentice" generations of paleontologists subsequent to the development of these practices would need to have inherited them from earlier "veteran" practitioners. And third, the magnitude of differences between versions of a practice should track the magnitude of differences in the fitnesses of those versions (*continuity*); smaller differences between practices should translate into smaller contributions to fitness, and likewise for bigger differences.

I hope to show over the course of this discussion that extinction satisfies all these criteria, and that this satisfaction explains why the general recognition of paleontology as an autonomous discipline—a "generator of testable and irreducible theories about pattern in the history of life"—coincides with the rise of extinction as a principal focus of paleontological research. Fueled by data from the fossil record, the study of extinction under the species-as-particles problem-solving framework revealed novel and unintuitive features of evolutionary history that could not have been predicted or disclosed by population genetics. More importantly, this approach to extinction offers special insight into the evolutionary process in general. By arguing persuasively for the general significance of extinction to our understanding of how evolutionary systems work, paleontologists were able to take their rightful place at the high table from which they had "too long been missing" (Maynard Smith 1984, 402).

9.1 Darwin Does Extinction

> The whole subject of the extinction of species has been involved in the most gratuitous mystery.
>
> —Charles Darwin, *On the Origin of Species* (1859, 318)

In his wonderful primer *Extinction: Bad Genes or Bad Luck?*, Raup (1991) points out that extinction has been bizarrely neglected as a scientific problem for which we ought to seek a theoretical understanding:

> If 99.9 percent of all species that have lived on earth are extinct, it follows that the total of species originations has been virtually the same as the total of species extinctions. Although present biodiversity—the millions of living species—seems high to us, today's biota results from a minor surplus of speciations over extinctions, accumulated over a long time. In view of these figures, it is puzzling that even evolutionary biologists have devoted almost no attention to extinction. Large monographs and textbooks have been written about speciation, and careers

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> have been built around the subject. But extinction has barely been touched. It's a little like a demographer trying to study population growth without considering death rates. Or an accountant interested in credits but not debits. Textbooks of evolutionary biology contain little about extinction beyond a few platitudes and tautologies like "Species go extinct when they are unable to cope with change" or "Extinction is likely when population size approaches zero." *The Encyclopedia Britannica* (1987) says, "Extinction occurs when a species can no longer reproduce at replacement levels." These statements are almost free of content. (11–12)

This was not always the case. In particular, for Darwin, extinction played a critical role at different stages of his evolutionary thought. As a way of introducing extinction as a theoretical concept, I will give a brief overview of some of the important ways in which extinction figured in Darwin's theorizing. My hope is that, in addition to illustrating a number of conceptual links between Darwin's treatment of extinction and contemporary treatments, this overview will lure the reader away from thinking of extinction merely as what happens when the last member of a species dies to thinking of extinction rather as a force that affects the evolutionary process at different levels of biological organization.

9.1.1 Extinction and the Harmony of Nature

> I think Case must be that one generation then should have as many living as now. To do this and to have many species in same genus (as is) *requires* extinction.
> —Charles Darwin, Notebook B (1837–1838, 36)

During Darwin's time there was a widespread metaphysical commitment to some kind of balance of nature or harmony of nature as a fundamental constraint on what the world must be like. For some, such as Lyell and the early Darwin, this principle was manifested in the geological record, which seemed to show that when one species goes extinct, a species that resembles it very closely seems to always follow immediately in space and time (the "law of succession of types"). This process of replacement was related to the notion, common among adherents of natural theology, that the number of species is held constant across time by an equilibrium process with strong parallels to the MacArthur-Wilson migration-extinction equilibrium described in chapter 8:

> If the ~~gradual deaths~~ existence of species is allowed each according to its kind, we must suppose deaths to follow at ['one after'] different epochs, & then successive births must repeople the globe or the number of its inhabitants has varied

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> exceedingly at different periods.—A supposition ~~fact~~ in contradiction to the fitness which the Author of Nature has now established. (Darwin 1835, quoted in Kohn 1980, 70)

Darwin followed Lyell in thinking that Nature's harmony required that species number be kept constant, but the effect on his early theorizing was often one of disharmony. Prior to reading Malthus in September 1838, Darwin's theory of how adaptation occurs was a version of Lamarckian "use and disuse," with a particular emphasis on his (Darwin's) view that new, adaptive *structures* arise as a direct result of the adaptive *habits* of preceding generations: "We even see they [actions] must be done often to be habitual or of great importance to cause long memory.—structure is only gained slowly.—therefore it can only be those actions, which many successive generations are impelled to do in same way" (Darwin 1838, 171, quoted in Kohn 1980, 132). Unlike his (and our) later theory, Darwin (like Lamarck) understood adaptation to be automatic, in the sense that all variation was adaptive variation. Variation, he believed, was the result of changes in the reproductive system that occur only to accommodate changes in the environment (Ospovat 1983, 41–43).

Now in Lamarck the adaptive process is one in which a species changes into another, better adapted species to meet the new demands of a changing environment. In this way, Lamarck was able to observe the Harmony of Nature constraint on species number while also avoiding a commitment to extinction, which he regarded as theologically unacceptable and whose circumvention was the inspiration for his theory of species transmutation (Burkhardt 1977, 130). Lyell, writing in 1830 and without the luxury of denying the reality of extinction, held the opposite but equally harmonious view: that new, perfectly adapted, and immutable species are created ex nihilo after other, formerly perfect immutable species gradually die as a result of their essential inability to adapt to a gradually changing environment.

For Darwin, adaptive transmutation occurs by sequestering adaptive variation until such time as it becomes fixed in heredity (Kohn 1980, 132; Ospovat 1983, 44–51). This process forms the early basis of Darwin's image of a branching phylogeny, where existing species produce new species as offshoots of a parent population. This is where things become difficult for Darwin. In contrast to Lamarck, who can maintain the balance of nature

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

by describing a process of adaptive transmutation that simply changes but does not add species, and Lyell, for whom balance is maintained by divine fiat, the mechanism envisioned by Darwin results in the multiplication of species and thus requires extinction to maintain balance:

> With this tendency to change (and to multiplication when isolated) *requires deaths of species to keep numbers of forms equable*. But is there any reason for supposing number of forms equable: This being due to subdivisions and amount of differences, so forms would be about equally numerous. . . . Organized beings represent a tree, *irregularly branched*; some branches far more branched,—hence genera.—As many terminal buds dying, as new ones generated. There is nothing stranger in death of species, than individuals. (Darwin 1837–1838, 21, quoted in Ospovat 1983, 55; emphasis added)

That Darwin needed extinction is clear; that he could have it is less so. After all, if adaptation is automatic, and if the environment changes at a responsible Lyellian rate, what on Earth could propel a species to extinction? One scenario he envisioned was where some new adaptive variation had been around long enough to become reasonably widespread but not long enough to have become indelibly fixed in the heredity of the new species. When the new, alternately adapted community rejoins the parent population, the adaptive variation is eventually blended out of the population, thus causing the "destruction" of the new species (Ospovat 1983, 56–57). It's worth noting that this scenario does not differ in its essentials from Futuyma's (1987) widely accepted explanation for the fossil pattern known as "punctuated equilibrium" discussed in chapter 8.

To kill well-established species, Darwin would employ one of the extinction devices formulated by Lyell. Although Lyell was committed to gradual change, he was not above allowing that slow, small environmental changes can trigger locally catastrophic events. Similarly, Darwin imagined that such local catastrophes would outpace an automatic adaptation process grounded in hereditary forces that were famously slow to change:

> Changes in structure being necessarily excessively slow they become firmly embedded in the constitution, which [is an-] other marked difference in the varieties made by . . . nature & man.—The constitution being hereditary & fixed, certain physical changes at last become unfit, the animal cannot change quick enough & perishes.—Lyell has shown such Physical changes will be unequally rapid with respect to their effects. (Darwin 1838, 153, quoted in Ospovat 1983, 57)

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

One of the tasks Darwin set for himself early in his theorizing was to explain why the natural system of classification was a nested hierarchy of groups subordinate to groups—that is, why certain species are naturally ranked as belonging to the same genus, certain genera are naturally ranked as belonging to the same family, and so on. By 1844 he believed that this grouping pattern followed as natural consequence of a branching phylogeny in which some species go extinct.

To see how this might work, imagine a world in which immortal species multiply in a branching process and offspring species differ only slightly from their parents' form (as per Darwin's insistence on only slight variations). This world's species would not form any natural groups, because "all forms would be separated only by small gradations" (Ospovat 1983, 172). In other words, for any three species *x, y, z*, the degree of morphological disparity between any pair of species [*x, y*] is equivalent to any other pair: [*x, y*] = [*x, z*] = [*y, z*]. If we were to take some species out of that world, it would no longer be the case that degrees of disparity were equivalent across all pairs. Rather, some species would be more similar to each other than to other forms, and thus some pairs would have lower disparity than other pairs. Sets of species that have relatively low disparity constitute a genus, and sets of genera that have relatively low disparity constitute a family, and so on up the hierarchy. Without differences in degrees of disparity between species pairs, there would be no reason to group *x* with *y* rather than with *z*, and thus no natural groups would exist. Extinction is what accounts for the natural groupings, because extinction is what generates the differences in disparity (see figure 9.1). In Darwin's words, "A genus is a real thing, but is consequent solely on extinction, if we take all organisms which have lived" (Darwin 1845, 40v).

Darwin would eventually abandon this picture as a sufficient explanation for why living things formed nested hierarchies. By November 1854, Darwin had adopted the view that there was a natural "tendency to diverge" in all taxonomic groups (Kohn 2008, 103). Because it makes disparity more apparent than real, the picture described above would have been consistent with a highly circumscribed morphospace, a commitment that Darwin could no longer have countenanced once he believed that a tendency for species to diverge morphologically was a law of nature (Ospovat 1983, 176). Thus at least part of the explanation for why living things exhibit such dramatic morphological differences is because, independently of extinction,

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

living things in general actually do differ quite a bit. However, as is well known, Darwin would still retain some connection between extinction and classification all the way through the publication of the *Origin*:

> extinction . . . has played an important part in defining and widening the intervals between the several groups in each class. We may thus account even for the distinctness of whole classes from each other—for instance, of birds from other vertebrate animals—by the belief that many ancient forms of life have been utterly lost, through which the early progenitors of birds were formerly connected with the early progenitors of the other vertebrate classes. (Darwin 1859, 431)

9.2 Extinction in the *Origin*

> Natural selection . . . leads to divergence of character and to much extinction of the less improved and intermediate forms of life. On these principles, I believe, the nature of the affinities of all organic beings may be explained.
>
> —Charles Darwin, *On the Origin of Species* (1859, 119)

The explanation developed to explain the "universal" "tendency to diverge" was what Darwin called the "Principle of Divergence," and it plays an immensely important role in the *Origin*. By September 1856, Darwin had become convinced that biological diversity itself, and the consequent natural system of classification, was due to natural selection favoring highly variable taxonomic groups—species with many varieties, genera with many species, and so on. His favorite model for this process was that of a little plot of land that can support some fixed amount of life. Imagine that two closely related species, *P* and *Q,* are the initial inhabitants of that plot of land, and that *P* is wide-ranging and contains many varieties, while *Q* is short-ranging and contains few varieties. Because *P* is more diverse, it is more likely than *Q* to produce a variety that can exploit some resource in that plot of land that neither *P* nor *Q* have evolved to consume: "The more diversified the descendants from any one species become in structure, constitution, and habits, by so much will they be better enabled to seize on many and widely diversified places in the polity of nature, and so be enabled to increase in numbers" (Darwin 1859, 112, quoted in Kohn 2008, 88).

Call *P*'s differently abled descendant *R*. Since members of *R* have unfettered access to otherwise unwanted resources, they will not struggle against members of *P* or *Q* to make a living on the little plot. But members of *R* will struggle against each other, and the ensuing intraspecific battle will favor

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

both (1) members of *R* who can improve their exploitation of *R*'s pet resource (standard directional selection), and (2) members of *R* who can find other, currently unwanted resources to exploit—what Darwin described as a "division of labor." These differently abled members of *R* will themselves struggle against one another for access to the new resource, initiating another directional selection regime that favors individuals who can improve their exploitation of that resource. Meanwhile, *P* continues to produce other, differently abled variants, some of which will be capable of exploiting yet further untapped resources, others of which will be capable of outcompeting *Q* for its own pet resources, driving *Q* to extinction: "After the foregoing discussion, . . . we may, I think, assume that the modified descendants of any one species will succeed by so much better as they become more diversified in structure, and are thus enabled to encroach on places occupied by other beings" (Darwin 1859, 116, quoted in Kohn 2008, 94).

There are three effects of these parallel process that are important for our purposes here, one concerning classification, one concerning extinction, and one concerning adaptation. The first is the particular form of growth exhibited by the evolutionary group *P*. *P* starts off as a highly variable species among many species in a genus, one of which was *Q*. *P*'s higher propensity to vary results in a higher propensity to speciate. And *P*'s higher propensity to speciate results in the eventual extermination of *Q*. Because of its higher propensity to speciate, *P* will someday grow to be a genus with many species and *Q* will meet a painful death at the paws of one of *P*'s many descendants. This growth process, rather than extinction, ends up being, in Darwin's view, the most significant factor in explaining why taxa form nested hierarchies. Given that some species actually are more similar to each other than they are to other species, in the same way that siblings are more similar to each other than to their cousins, and that some species have an extraordinary propensity to speciate while others are far less fecund, the actual history of life will "represent a tree, *irregularly branched*; some branches far more branched,—hence genera.—As many terminal buds dying, as new ones generated" (Darwin 1837–1838, 21).

Darwin's discussion of the Principle of Divergence in chapter 4 of the *Origin* also offers insight into what he saw as the chief mechanism by which species are extinguished, competition for resources. Each species is perennially beleaguered by competition from better adapted varieties, daughter species, and more distantly related species for a limited pool of

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

resources. Intense competition within individual lineages serves to adaptively "improve" members of that lineage, pushing them toward "perfection" and exterminating "intermediate forms," while competition between lineages similarly raises the aggregate adaptedness of a genus by eliminating less fit species:

> The modified offspring from the later and more highly improved branches in the lines of descent, will, it is probable, often take the place of, and so destroy, the earlier and less improved branches: this is represented in the diagram by some of the lower branches not reaching to the upper horizontal lines. (Darwin 1859, 119, quoted in Kohn 2008, 100)

This picture of how species are killed is helpful to Darwin not only because it solves the problem of classification that he had set for himself, but in particular because of (1) the phenomena it allows him to explain and (2) what it implies about evolutionary progress. Regarding (1), the absence of intermediate forms either living today or in the fossil record—one of the "difficulties on theory" identified by Darwin in the *Origin* in chapter 6 and again in chapter 9—are both ultimately explained by the rate at which lesser adapted forms are killed off on the inexorable death march to "extreme perfection." Darwin imagines that a population of intermediate forms will exist for a relatively short time and be kept relatively small owing to their futile struggle against the always present, slightly better adapted members of the same lineage, and that these populations will remain small until the lineage is nearly perfect, at which time stabilizing selection takes over and maintains the population in a more or less morphologically consistent state—"for in all cases the new and improved forms of life will tend to supplant the old and unimproved forms" (Darwin 1859, 281). Similar considerations apply to another well-known feature of the fossil record, what he calls "the Forms of Life changing almost simultaneously throughout the World" (322):

> As it seems to me, the parallel, and, taken in a large sense, simultaneous, succession of the same forms of life throughout the world, accords well with the principle of new species having been formed by dominant species spreading widely and varying; the new species thus produced being themselves dominant owing to inheritance, and to having already had some advantage over their parents or over other species; these again spreading, varying, and producing new species. The forms which are beaten and which yield their places to the new and victorious forms, will generally be allied in groups, from inheriting some inferiority in common; and therefore as new and improved groups spread throughout the world,

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> old groups will disappear from the world; and the succession of forms in both ways will everywhere tend to correspond. (Darwin 1859, 327)

In both cases, the fact that intermediate forms are small and short-lived means that the probability of any intermediate form fossilizing is very small, which explains why they are so rarely discovered. Here again we see a strong parallel to Futuyma's (1987) argument that most evolutionary change is ephemeral, which explains its absence in the fossil record and the consequent apparent stasis (see figures 9.1 and 9.2).

One final feature of Darwin's writing on extinction is worth mentioning because of how it connects with contemporary theory. Section 9.1 begins with a quotation from the *Origin*, where Darwin expresses some frustration over why anyone (including himself) would have ever been bewildered by extinction ("No one I think can have marveled more at the extinction of species, than I have done" [Darwin 1859, 318]). To understand extinction, he says, we can begin by noting the general fact that "rarity precedes extinction" (319). Of course, rarity never surprised anyone, and indeed is very easy to understand, as Darwin illustrates by relating a story about a fossil tooth from an extinct horse:

> Had this horse been still living, but in some degree rare, no naturalist would have felt the least surprise at its rarity; for rarity is the attribute of a vast number of species of all classes, in all countries. If we ask ourselves why this or that species is rare, we answer that something is unfavourable in its conditions of life. (319)

There's no mystery here: "the increase of every living being is constantly being checked by unperceived injurious agencies" (Darwin 1859, 319); that any particular species is being held down is just more of the same. And extinction is but a stone's throw from rarity: "If the conditions had gone on, however slowly, becoming less and less favourable, we assuredly should not have perceived the fact, yet the fossil horse would certainly have become rarer and rarer, and finally extinct;—its place being seized on by some more successful competitor" (319).

What is distinctive about his discussion of rarity is the way in which he suggests that it is simply obvious that a species' rarity is caused by the same "injurious agencies" that would eventually kill it off. Even in cases of "the *apparently* sudden extermination of whole families or orders" (emphasis added), Darwin uses the same explanatory strategy to demonstrate that the instantaneousness is merely apparent:

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

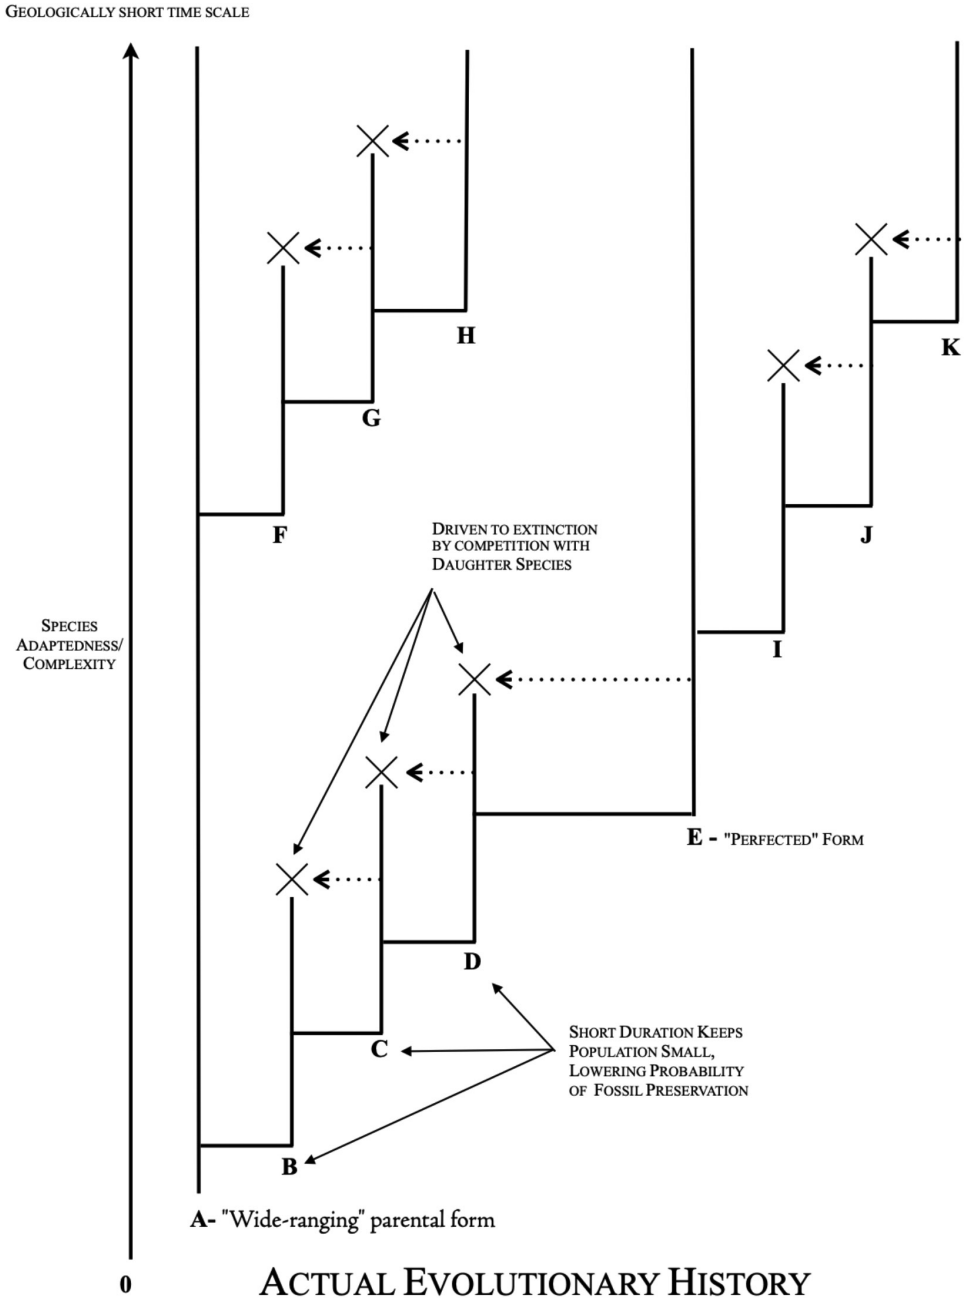

Figure 9.1

Darwin's model of the evolutionary process that results in "gaps" in morphology. Letters indicate names of hypothetical lineages; crosses indicate extinction of a lineage; dotted lines indicate one lineage driving another to extinction.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

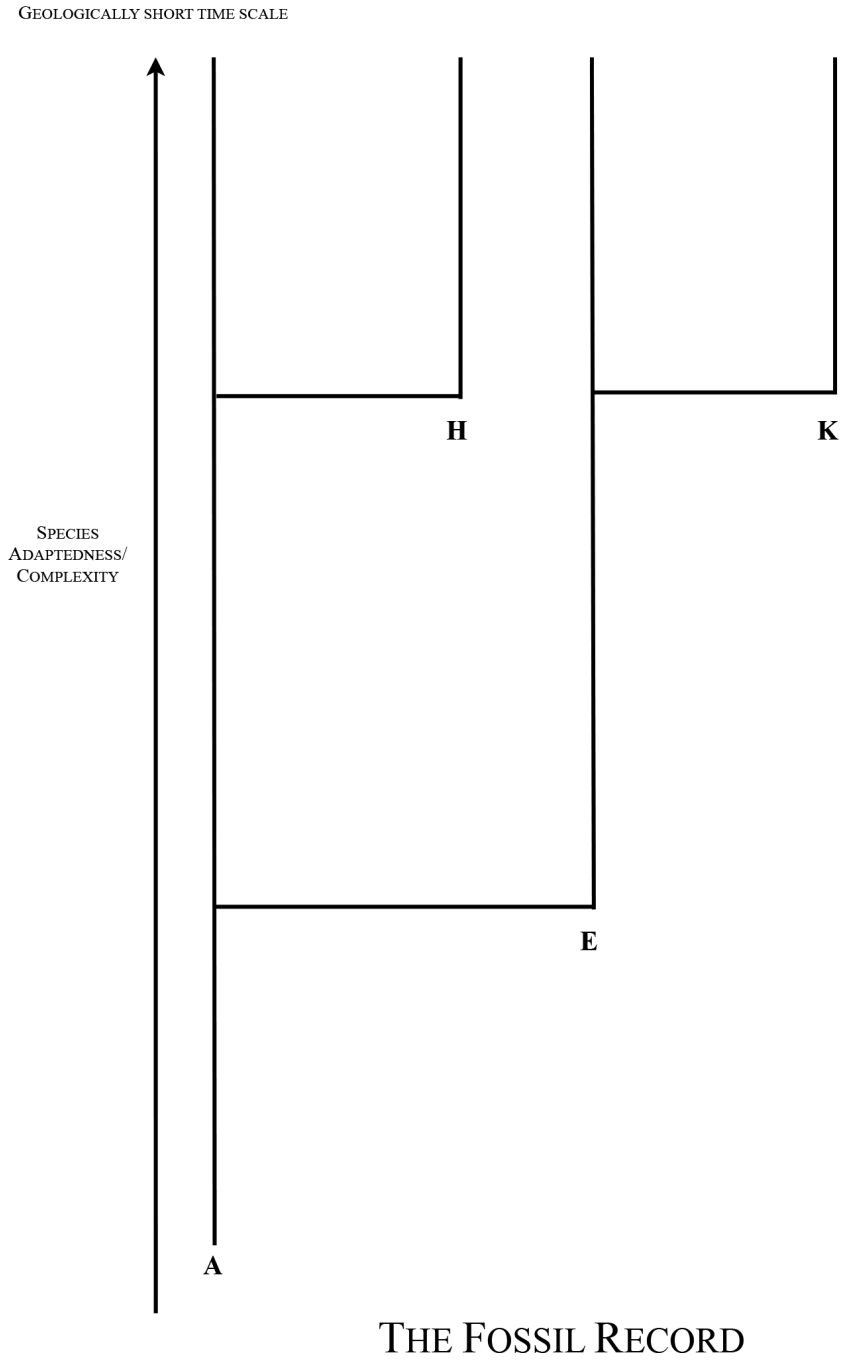

Figure 9.2
The fossil record, based on the process depicted in figure 9.1.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> With respect to the apparently sudden extermination of whole families or orders, as of Trilobites at the close of the palæozoic period and of Ammonites at the close of the secondary period, we must remember what has been already said on the probable wide intervals of time between our consecutive formations; and in these intervals there may have been much slow extermination. Moreover, when by sudden immigration or by unusually rapid development, many species of a new group have taken possession of a new area, they will have exterminated in a correspondingly rapid manner many of the old inhabitants; and the forms which thus yield their places will commonly be allied, for they will partake of some inferiority in common. (Darwin 1859, 321–322)

First of all, he urges, let's remember that what looks rapid in the geological column ain't exactly greased lightning. Having said that, we can no doubt imagine conditions under which the rate of competitive extermination is accelerated, as when an immigrant group of closely related species wreaks havoc on the ill-prepared denizens of their new homeland. But this process of extermination does not differ in kind from the one in which a horse species is gradually outcompeted over several million years. In both cases these evolutionary groups have entered a losing battle against "injurious agencies"; it's just that some battles end more quickly than others.

Whether one is inclined to agree with Darwin or not, the power and simplicity of his reasoning cannot be denied. The sheer breadth of phenomena that he is able to derive from the deeply conservative assumptions that organisms vary, that they often pass their variations on to their offspring, and that these variations cause differences in reproductive success, is nothing short of humbling. Extinction is a case in point: "The theory of natural selection is grounded on the belief that each new variety, and ultimately each new species, is produced and maintained by having some advantage over those with which it comes into competition; and the consequent extinction of less-favoured forms almost inevitably follows" (Darwin 1859, 320). From this seemingly trivial point follows an impressive array of widely accepted generalizations, such as that organisms can be classified in groups subordinate to groups, that these groups will be many and varied, that extant forms will be nearly perfectly adapted, and that the fossil record will be unlikely to retain transitional forms. With equal facility, he is able to accommodate as recalcitrant a datum as the abrupt disappearance of higher taxonomic groups.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

9.3 Modern Perspective on Extinction

> I am becoming more and more convinced that the key gap in our thinking for the last 125 years is the nature of extinction.
>
> —Raup to Schopf, January 28, 1979, Smithsonian Institution Archives, Thomas J. M. Schopf Papers, Washington, DC

> I find it surprising: I would have expected a major cause of extinction to be competition from other taxa.
>
> —John Maynard Smith, "Paleontology at the High Table" (1984, 402)

Considering the simplicity and intuitiveness of the competitive extermination model of extinction, Maynard Smith's expectation is entirely warranted. But the 1980 discovery of unnaturally high levels of iridium at the precise point in the stratigraphic column where the dinosaurs disappear suggested that their extinction at the end of the Cretaceous period (known as the "K-Pg extinction") had occurred relatively quickly as the result of a catastrophic bolide impact (the only plausible source of the elevated iridium level), not gradually as the result of competition with better adapted forms. If competition with other taxa was plausibly not a major cause of extinction at the end of the Cretaceous—a case in which the majority of all species on the planet were killed off—then competition's monopoly on extinction is in serious jeopardy, as are many of the broader evolutionary consequences traditionally associated with the competitive mode of extinction. The Marine Biological Laboratory (MBL) project (see chapter 8) had similarly shown that nonselective extinctions could generate many of the morphological trends typically attributed to selective extermination of lesser adapted forms (particularly Raup and Gould 1974). After a successful and uncontested post-Darwinian run, the competitive model finally had some competition of its own.

But there was much more at stake than what killed the dinosaurs. The generality of the competitive model, combined with the generality and significance of the evolutionary consequences described by that model, mean that threats to its hegemony have the potential to fundamentally alter not only our understanding of our actual evolutionary history but also the very nature of the evolutionary process. In the wake of the MBL project, the handful of collaborators would begin to question the competitive model's primacy as an explanation for extinction as well as the host of evolutionary

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

trends that were supposed to fall neatly within its purview. This interrogation would eventually lead to a number of deeply illuminating and distinctively paleontological contributions to our understanding of the nature and evolutionary significance of extinction, both generally and with respect to Earth's unique history of life.

In this section I review some of the central alternatives to the competitive model of extinction and discuss the broader implications of these alternatives for our understanding of evolutionary history and the evolutionary process. I then go on to explain, via the philosophical model of disciplinary autonomy, how these alternatives and their implications form the conceptual basis for paleontology's emergence as an autonomous theoretical discipline.

9.3.1 In Search of a Null Model

It is fair to say that David M. Raup has done more to fill the gap in our thinking about extinction than anyone else in the history of biological thought. Using a species-as-particles approach that ignores "the peculiarities of individual taxa," Raup initiated a research program aimed at articulating alternatives to the standard competitive extermination model, thereby encouraging a radical reassessment of the nature of extinction. These alternatives take as their starting point the question of whether, in Raup's words, "proneness to extinction is an inherent property of a species—a weakness—or does it depend on the vagaries of chance in a risk-ridden world" (1991, 5). In other words, do "the peculiarities of individual taxa"—their specific phenotypic properties—affect the probability that they will go extinct, as Darwin thought it did? Or do all species exist in a "field of bullets," where species death is accidental, where who lives and who dies is a matter of chance? The intuitive answer, the one for which Darwin argued persuasively and which reigned supreme for 125 years, is that of course, the peculiarities of individual taxa matter; their specific vulnerabilities to certain "injurious agencies" will eventually result in their death, just as, Darwin said, a sick man's disease will eventually finish him off. But as Raup began to point out in various places in the late 1970s:

> If we take neo-darwinian theory at face value, the fossil record makes no sense. That is, if we have (a) adaptation through natural selection and/or species selection and (b) extinction through competitive replacement or displacement, then we ought to see a variety of features in the fossil record that we do not such as: (a) clear evidence of progress, (b) decrease in evolutionary rates (both morphologic

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> and taxonomic), (c) possibly a decrease in diversity (at least within an adaptive zone). Now we do not see these things because: (a) we are too dumb, *or* (b) the record is lousy, *or* there are features of the evolutionary mechanism that prevent the approach to a steady-state. (Raup to Schopf, January 28, 1979, quoted in D. Sepkoski 2012, 259–260; see also Raup 1978)

Using the conceptual framework developed in the MBL project, Raup would spend the rest of his professional life formulating and testing models of evolutionary processes that could both result in extinction and make sense of the fossil record. At the core of these models is an assumption, contra Darwin, that there is an important element of randomness to the extinction process—that it was often merely a result of a run of "bad luck" rather than the predictable result of a species' losing battle with superior beings.

There are a few independent considerations that make it reasonable to assume that extinction is effectively random in certain respects. One is that patterns of extinction in the fossil record resemble what we would expect if species were randomly going extinct (this was one of the messages from the MBL project). Another is the fact mentioned in the Raup (1991) quote cited above, that 99.9% of all species that have ever existed have gone extinct. If that's true, then the long-term probabilities of extinction for all species are basically equal, which means that one species is no more likely to go extinct than another (Raup 1991, 53). Yet another reason derives from Van Valen's "Law of Extinction" (see chapter 8), which purports to show that extinction of any sized taxonomic group is random with respect to taxon age.

The assumption that extinction is random in some important sense allows us to model extinction on random processes that are well understood. One such process is casino gambling, where randomness plays a central role and intuitive role (see Raup 1991, chap. 3). Imagine a gambler with $10 who has sat down at a casino table to play a game where each bet is $1 and has 50/50 odds. One thing we know about what is going to happen is that this gambler will eventually lose all of his money. He may be able to stay in the game rather unimpressively for a while, winning a couple here and losing a couple there, but sooner or later his losses will catch up with him and he'll be out of the game. Or he may stay in the game rather impressively with a long and spectacular winning streak; even then, though, eventually he'll still lose all his money.

Why might not the gambler break the casino instead? This is because the casino gambling process described consists of a random walk with a very

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

distant upper limit and a very near lower limit. The lower limit is the point at which the gambler loses all his money and is unceremoniously ejected from the casino, and is therefore an absorbing boundary. The upper limit is the total amount of money available to the casino. It is also an absorbing boundary in the sense that the random walk stops if it hits this boundary—that is, if the gambler wins all of the casino's money. Now, the gambler will definitely reach one of these boundaries at some point; which one he reaches depends on which one he is closer to at the outset. Since mere mortal gamblers have coffers that pale in comparison to those of a casino, a gambler is guaranteed to lose all his money if he plays long enough; he can delay his fate by bringing more money to the table, but he cannot escape it so long as he starts off with less money than the casino (in real life this outcome is accelerated by casinos' preference for games with odds skewed in their favor).

We can model the extinction of a genus as that of a gambler in the game of life, where the genus's constituent species are its betting chips, and extinction is an absorbing boundary comparable to going broke—once the genus loses all of its species, it's out of the game. Rather than thinking of a particular genus as slated for extinction because of the peculiarities of its constituent species, imagine instead that each of its species (along with the species of all other genera) has 50/50 chance of going extinct in the next million years and a 50/50 chance of speciating during that same interval. In effect, each million years is a bet made by the genus that it will avoid extinction and survive until the next round.

Like the casino gambler, we do not need to know anything about its relative advantages, its clever strategies, or its ability to accommodate novel events in order to know that this genus is on its way out. It may take a little time, but eventually it will lose all of its species to the absorbing boundary of extinction, just like the gambler never fails to go broke. A genus can increase its time to extinction by speciating aggressively for many millions of years, but ultimately it cannot avoid its inevitable departure. That it might experience an acceleration in the rate at which it speciates does not depend on whether it evolves some highly adaptive novelty or other. Rather, it can do this for the same reason (or lack thereof) that a gambler might have a winning streak: merely a bit of good luck in an otherwise pitiless game. We also know, based on the similarity with the gambling process,

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

that the genus's number of constituent species may vary wildly over the course of its life, sometimes rising above—even high above—its initial frequency, at other times coming within a hair's breadth of annihilation, only to be brought back from the brink for another roll of the dice. That is, it is not on a steady decline from day 1. But it is destined for extinction in any case.

We can use this model to make predictions about what sorts of patterns we should see in the fossil record if species are randomly going extinct. These predictions tell how probable it is that a particular higher-order taxon will go extinct over a given interval of time purely by chance, akin to telling us the probability that a casino gambler will go broke after some specified number of bets. When we then turn to the fossil record for some taxon, we can use these predictions to assess whether its actual extinction rates over a specific time in evolutionary history are consistent with what we would expect if extinction was random. If they are not consistent (i.e., if they are either higher or lower), then we can start to ask more focused questions about what special cause could explain the deviation from chance expectations.

It is at this point that the peculiarities of individual taxa become relevant. For example, if we would expect a genus with 10 species to go completely extinct over 40 million years, and the fossil record indicates that a few of that genus's species lasted for, say, 300 million years, it then becomes appropriate to consider whether there was something about the peculiarities of those species that allowed them to beat the odds and keep their genus in the game for a while longer. Only then does an appeal to competitive advantage become warranted, because it is only then that we ought to be surprised that the genus went extinct when it did. This points to one major gap in Darwin's perspective on the fossil record. Because he had no sense of or interest in how the history of an evolutionary group would look in the absence of selective forces, he was unable to make a distinction between cases where an evolutionary group could have meandered to extinction merely by chance, and cases where an evolutionary group was driven to extinction by "injurious agencies." To Darwin, no extinctions were surprising, because natural selection is constantly improving the fitness of lineages, and extermination is the default expectation for inferior beings. But as the gambler analogy shows, extermination is also the default expectation for superior beings.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

9.3.2 The Null Model Rejected

The Gambler's Ruin model departs from Darwin's view in describing what sorts of extinction patterns to expect if the peculiarities of individual taxa do not matter. But this is not the only way to differ from Darwin. Recall that the model of extinction presented by Darwin in the *Origin* is one in which progressively more adapted forms drive their parent forms to extinction as the lineage approaches perfection, and that this process occurs both within and between evolutionary groups. The particular process described by Darwin is regarded as progressive because it favors organisms that outcompete others with respect to achieving some goal made relevant by the "conditions of existence"—that is, by the normal environment to which they are adapted. As long as nature doesn't move the goal posts too quickly, a lineage can by this process improve its degree of adaptedness in its environment. On this view, the peculiarities of individual taxa are highly significant, because it is a taxon's peculiarities that determine the degree of fit with its environment, which in turn determines the probability that the taxon will be represented in the next generation. As nature patiently favors some peculiarities over others, eventually all but the most perfect will be filtered out of the population. Unlike in the casino, where nothing inherent to the gambler makes a difference to his success, here the peculiarities of individual taxa matter—they matter big time.

We can imagine another scenario under which peculiarities matter, but not because of their degree of fit with the environment to which the taxon is adapted. This would be a case where the environment changes radically over a short period of time, thus changing, perhaps equally radically, which peculiarities are favored in that area. The shift in which peculiarities are favored will be associated with a corresponding shift in the direction in which adaptive progress proceeds. Whereas in times prior, natural selection was pushing a particular taxon toward more and more improved versions of one peculiarity, after the environmental perturbation this same taxon will find itself being driven toward better versions of some other peculiarity. Because the environmental change occurs so rapidly, those previously at the top of the heap might find suddenly find themselves hopelessly maladapted to the new conditions, with extinction right around the corner. As in the case of Gambler's Ruin, the relevance of competition to extinction here is marginal at best. Those organisms that are caught off-guard by

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

catastrophic environmental change would be ill-fated even with infinite resources available to them.

As an intuitive model of this process, consider my old school, where students are ranked and where the lowest ranked 50% are prevented from graduating to the next grade. In addition to being a bit cutthroat, this school was unusual in that it changed its ranking criteria every two years, a tradition that I only learned about long after having graduated. The schedule, as it turns out, breaks down like this:

Kindergarten: grades

First grade: family income

Second grade: height

Third grade: number of vowels in last name

Fourth grade: number of consonants in first name

Fifth grade: top speed

Sixth grade: how many push-ups you can do

Seventh grade: ratio of hopes to dreams

Eighth grade: shoe size

Ninth grade: weight

Tenth grade: top speed

Eleventh grade: number of hair follicles

Twelfth grade: how many push-ups you can imagine doing

My school's way of doing things had a number of interesting effects. One of them was that the group of students who survived the abrupt change of rules and went on to graduate typically had very little in common with the previously successful group. For example, the average first-grader was smarter than the average kindergartener, but the average second-grader was dumber than the average first-grader; by the time we got to seventh grade, intelligence was all but randomly distributed across the rankings. Or consider that the average second-grader was richer than the average first-grader, but the average third-grader was poorer than the average second-grader. One ninth-grader's income differed from another's solely by chance, while their average shoe size was predictably higher than that of the eighth-grade class. From fifth grade on, all students were from Eastern Europe.

Two features of this filtering process are worth mentioning. The first was that over the years, my class made relatively little progress along any of

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

these dimensions. Sure, only the smartest kids made it to first grade, but then the school stopped caring about brains, and the average intelligence of my class dwindled with each passing year. Similarly with family income and height. In our last year the school tried to turn things around with the bizarre idea that the diplomas should be given to those who could do the most push-ups, but push-up ability had been neglected for so long that the degree-earning group was rather unimpressive in its efforts to push up. Thus, although the school was highly selective in which students it would pass to the next grade, the rule-changing policy proved to be not very constructive.

Another feature concerned our ability to predict who would make it to the next class. Because the rules changed from year to year, we could not use the fact that certain students had made it into the next grade to predict who would be advancing in a year's time. Even though the school was being very careful in who it allowed to continue to the next level, it kept changing its criteria, making it impossible to know what the next class's constitution would look like. And longevity made no difference, since, obviously, this turned out to be true even in cases where students had made it all the way to grade 12: after avoiding expulsion for all those years, 50% of the students who had survived every cut were not allowed to graduate.

In biological terms, the expulsion of 50% of students at the end of the year amounts to a *mass extinction*, an event that, like the process described above, has some curious regularities associated with it. For example, we know that we cannot use our knowledge of past or present adaptedness to predict which taxa will survive a mass extinction event, something we learned from observing that:

> Factors such as local abundance, reproductive mode, body size and inferred generation time, trophic strategy, life habit, geographic range at the species level, and species richness, which have all been hypothesized or shown to be significant under ''normal'' extinction intensities, had little effect on genus survivorship during the K/T extinction and were unimportant in one or more of the other mass extinctions as well. (Jablonski 2005, 194)

Because of their radical restructuring of the biosphere, "these rare, intense events derail the evolutionary patterns shaped by organismic selection, speciation, and species sorting during times of background extinction" (192), with the result that the taxa most successful before a mass extinction tend to be far less so in the aftermath—that is, if they're still around. In the same way that my school favored different peculiarities from year to year, the

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

peculiarities of individual taxa that are favored in "background" times differ from those favored during mass extinction events, producing the same lack of progress along particular dimensions that was witnessed among my classmates. Raup called this feature of the evolutionary process "nonconstructive selectivity," and it is hypothesized to be a general feature of any evolutionary system exposed to perturbations of a sufficient magnitude: "This kind of extinction is selective but not constructive; it does not lead to organisms better able to survive in their normal environment . . . I call this kind of selective but nonconstructive extinction *wanton* extinction, on the basis of the word's original meaning: undisciplined or unruly" (Raup 1991, 187).

This is not to say that the survivors of mass extinctions are randomly determined. Since the early 1980s, paleontologist (and former Raup colleague) David Jablonski has argued persuasively for the principle that a clade's geographic range can tell us about its chances of surviving a mass extinction—call this principle "Jablonski's Law." Jablonski's Law states that a genus's having a broad geographic range increases its probability of surviving a mass extinction, regardless of the breadth of constituent species' geographic ranges. There is also evidence that the law is invariant across other kinds of peculiarities of constituent species. "For example, Late Cretaceous marine bivalve and gastropod genera that contained many, mainly widespread species tend to have significantly greater durations than genera consisting of few, spatially restricted species, and the other combinations tended to show intermediate values" (Jablonski 2005, 195). This trend is predicted by the Gambler's Ruin model discussed in section 9.3.1: when species are going extinct at random, we expect genera with more species to outlive genera with fewer species.

What is surprising is that the expectation that species-rich genera will outlast species-poor genera is violated during the mass extinction at the end of the Cretaceous period, where species-rich and species-poor bivalve genera show equivalent survival rates (Jablonski 1986, 132). Even more bizarre is the fate of gastropods, where species-poor genera fared far better than species-rich genera (47% vs. 29% surviving, respectively). Most important, though, is the fact that "geographic ranges of constituent *species* had no significant effect on clade survival, but geographic range of clades *did* have a significant effect" (131; emphasis added). This is like saying that the widespreadness of nuclear family members makes no difference to whether an extended family survives, while the widespreadness of nuclear families does make a difference to whether the extended family survives. The widespreadness of

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

nuclear families is not an artifact of the widespreadness of family members, in the sense that the widespreadness of families can vary independently of that of their constituent members: a wide-ranging extended family may have constituent nuclear families who live on different continents but whose members live in a single room; a narrow-ranging extended family may have constituent nuclear families whose members are randomly distributed across the extended family's entire range. To increase its chances of surviving a mass extinction, all that matters is that an extended family be wide-ranging. Once again, the peculiarities of individual members do not matter.

Because mass extinctions favor the breadth of a clade's geographic range, not the adaptedness of its constituent species, it is probable that "even well-established clades and adaptations could be lost during these episodes, simply because they were not associated with the features that enhanced survivorship during these unusual and geologically brief events" (Jablonski 2005, 197). The tendency of mass extinction events to deny exceedingly well-adapted species the preferential evolutionary treatment to which they have grown accustomed during background times has the effect of leveling the adaptive topography in a region, allowing for the "subsequent diversification of formerly marginal taxa" (197). It is to this subsequent diversification that we now turn.

9.4 The Evolutionary Significance of Extinction

> This argument can be made more concrete with simulations.
>
> —J. John Sepkoski Jr., "A Kinetic Model of Phanerozoic Taxonomic Diversity" (1978, 238)

Recall that Darwin believed natural selection to be the evolutionary mechanism most responsible for the diversification of life on Earth: evolutionary systems generate diversity by favoring the development of new ways to exploit unused resources. The process envisioned by Darwin emphasizes the significance of the peculiarities of individual taxa for allowing them to occupy niche space, where organisms' idiosyncracies promote divergence within and between lineages to the extent that those idiosyncracies complement some unoccupied niche or other. For Darwin, this process of diversification causes but does not require extinction. Nature promotes and then proceeds to improve divergent forms as long as those forms' peculiarities correspond to some available "places in the polity of nature"; extinction

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

is an incidental by-product of the fact that some forms' peculiarities leave them better suited for holding office.

By contrast, the modern view on diversification accords pride of place to extinction, not natural selection. And, in keeping with the dominant contemporary models of extinction, it does not place any particular importance on the peculiarities of individual taxa. Like Raup's work on extinction, modern views on diversification grew predominantly out of the conceptual (species-as-particles) and methodological (computer simulations) lessons associated with the MBL project.

The significance of the MBL work to contemporary models of diversity has its roots in the work of Gould's star graduate student, J. John "Jack" Sepkoski Jr. Coming under Gould's tutelage in 1970, Sepkoski played an important technical role in the original MBL studies and would eventually refine the MBL approach to the point where it could shed light on, as well as reveal, a number of key evolutionary problems. He would, in addition, use his refined version of the MBL model to argue for the conceptual coherence and methodological reliability of the species-as-particles problem-solving approach. In Jack Sepkoski and others of his generation (such as Jablonski), we see the realization of the remaining two conditions for qualifying as a Darwinian population: *heredity* and *continuity*. I will discuss each of these elements in turn, connecting them to the eventual emergence (or *re*emergence, in Gould's imagination) of paleontology as an autonomous scientific discipline—a Darwinian population of scientific practices.

In 1978, Sepkoski published the first of what would eventually be a three-part series of papers on the history of marine taxonomic diversity during the Phanerozoic, the geological time period stretching from the beginning of the Cambrian 542 million years ago to today—basically the entire history of life. Uniting these papers is what Sepkoski described as "a modification of the 'MBL' computer program developed by Raup et al. (1973)"—an exquisitely simple model designed to show what we should expect the history of diversity to look like if we assume that the process of "large-scale diversification over evolutionary time" is a MacArthur-Wilson style equilibrium process governing changes in origination and extinction rates (J. Sepkoski 1978, 224). As in the original MBL work, Sepkoski's model assumes nothing about individual taxa other than that they possess some probability of branching and some probability of going extinct during a given time interval. Like MacArthur-Wilson, the initial state of this model is one with no diversity (and thus no extinction) and a high per-taxon rate of

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

speciation. The system then proceeds to diversify up until the point where the rate of speciation equals the rate of extinction, at which time the system is in equilibrium (with any perturbation of speciation rates resulting in a commensurate response in extinction rates). In an effort to make the model more realistic, Sepkoski injected a stochastic element that allowed "the per species rates of origination and extinction to vary with slight random components, representing perturbations in the environment, peculiarities of the taxa involved, etc." (233; see figure 9.3). As he observed, "The entire Phanerozoic history of marine ordinal diversity is eminently comparable to this stochastic curve. . . . The logistic pattern is so striking that most specific

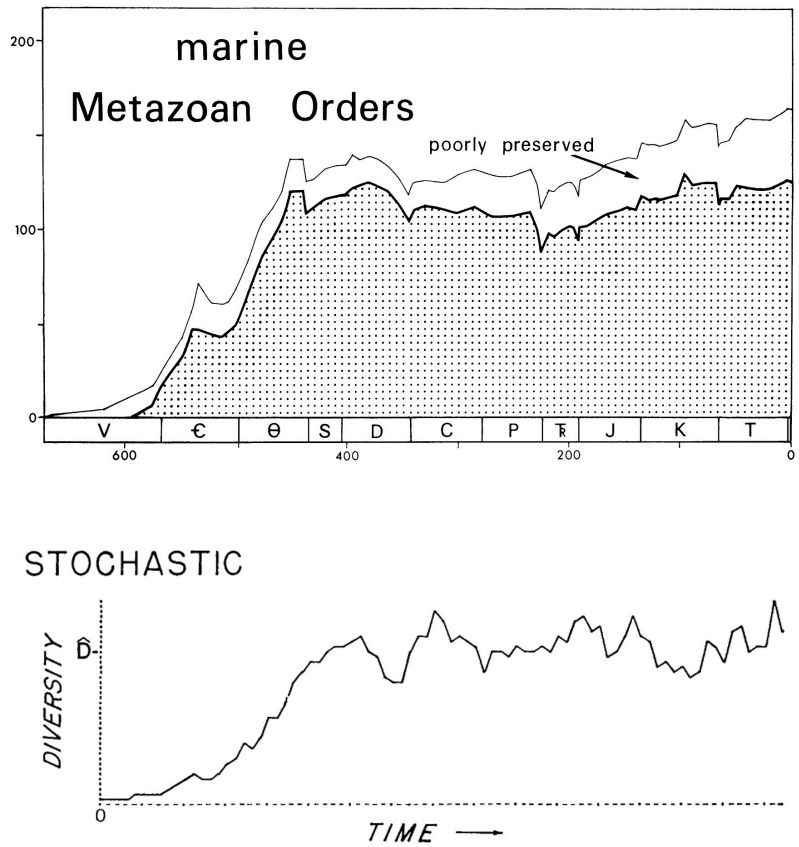

Figure 9.3
Top: Sepkoski's simulation of long-term diversity using a species-as-particles approach. *Bottom*: The actual Phanerozoic diversity pattern for marine orders (from J. Sepkoski 1978, 233–234).

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

historical events, represented by knicks and bumps in the curve, appear no more important than simple stochastic fluctuations" (233).

One important feature of the Phanerozoic that Sepkoski (1978) found his model unable to reproduce was "the chaotic pattern of ordinal extinctions," upon which basis he argued that "other factors, particularly periods of extraordinary extinction, must also have been operative" in producing the history of life entombed in the fossil record (244). In a 1984 follow-up paper (the last in the three-part series), Sepkoski introduced what he called "time-specific perturbations" designed to coincide with major mass extinction events known to have occurred during the Phanerozoic. Combined with an updated 1979 version of the model which divided life into three distinct "Phases," each characterized by a dominant evolutionary fauna (Cambrian [Cm], Paleozoic [Pz], and Modern [Md]), Sepkoski was able to simulate the large-scale pattern of diversification over the entire Phanerozoic with astonishing precision (figure 9.4).

Because of the gravity of the result and Sepkoski's palpable astonishment, it is worth quoting his reaction at length:

> Yes, the fossil record does appear highly structured, at least at the scale of changes in familial diversity within suites of higher taxa. Given that evolutionary faunas behave, to a first approximation, like discrete entities, much of the evolutionary change in diversity and faunal composition that we see in the Phanerozoic oceans seems to be a consequence of a very small number of parameters and their functional relationships. Most of the change can be described in terms of initial per-taxon diversification rates and apparent equilibria, which, in the three-phase kinetic model, are intrinsic and invariant properties of the three evolutionary faunas. This holds true even for the changes effected by mass extinctions; although the timing and intensity of these events appear to be controlled by forces external to the faunal system, the responses of the faunas during decline and rebound seem to be properties of their previously set initial diversification rates and equilibria. Therefore, much of what appears complex, or even random, in the history of familial diversity and faunal change may not originate from any inherent complexity but rather from the nonlinear functional interrelationships among the set parameters. (J. Sepkoski 1984, 262)

Figure 9.4
Top left: Actual Phanerozoic diversity for marine families. *Bottom left*: Sepkoski's simulation of long-term diversity with three interacting but distinct evolutionary faunas (1, 2, 3) and with successively more mass extinction events (A, B, C) (from J. Sepkoski 1984). The top figure is known as the Sepkoski Curve.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

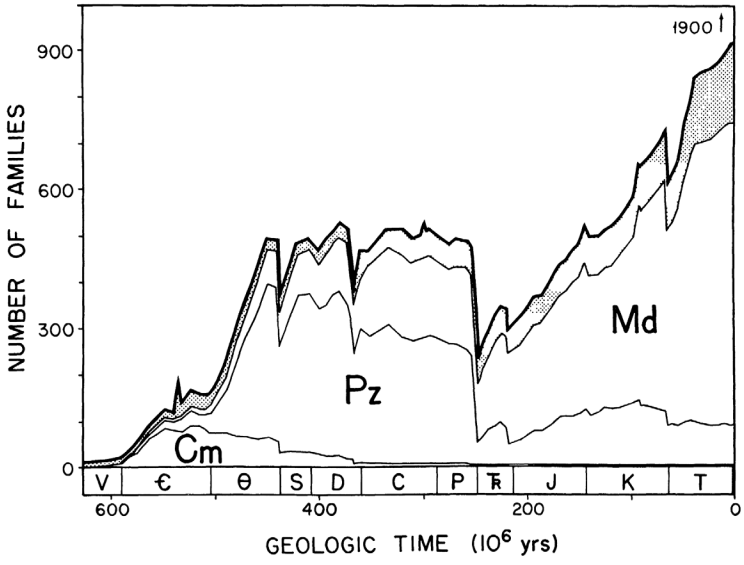

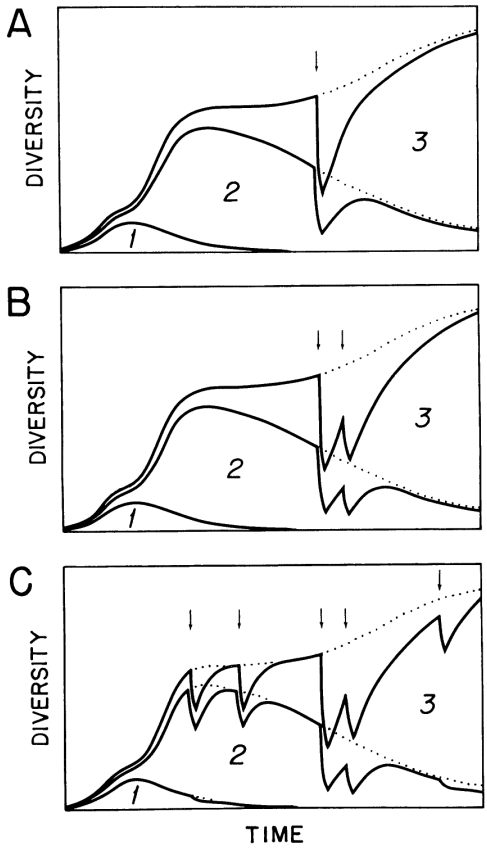

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

As Gould delighted in pointing out, the particular phenotypes produced over the course of evolutionary history could have—indeed, would have—looked completely different were we to replay the tape of life. What would not look different (at least not significantly different) is the pattern of biological diversification—a respectable sigmoidal growth pattern, variously peppered with random fluctuations around an equilibrium diversity. Even at a glance, there is no mistaking the fact that Sepkoski's model suggests that when it comes to generating the large-scale pattern of diversification reflected in the fossil record, the peculiarities of individual taxa do not matter. But then, what does matter?

What seems to matter is extinction—in particular, extinction of the wanton variety that characterizes mass extinction events. We know that the major adaptive radiations that have occurred in Earth's evolutionary history have followed the extinction of many higher-order taxa. The origination-extinction equilibrium process is part of the explanation of this trend: the extinction of lots of higher-order taxa seen during these times drives down diversity, causing a corresponding increase in origination rates (what is known as the "rebound" or "recovery" period). This model, combined with the actual fossil record, suggests that extinction rates of higher-order taxa need to rise in order to achieve the level of diversification witnessed during the historically important adaptive radiations.

We also know that the elimination of higher-order taxa in general is very improbable if species are going extinct purely by chance. If species are randomly killed off, irrespective of their higher-order group membership, then the probability of a higher-order groups surviving increases as we travel up the group hierarchy. For example, assuming that each species has a 50% chance of going extinct (frighteningly high, but a relatively modest rate by mass extinction standards), a genus with 10 species has about a 0.01% chance of going extinct, because each of its species would have to randomly die off in order to annihilate the genus; you can imagine how hard it would be to kill off a family, order, or class in this way.

This suggests that mass extinctions cannot be characterized by random species extinction. Rather, the scale of loss witnessed during these events requires that the odds be stacked against certain kinds of organisms. As Raup suggests,

> If some biological groups are really harder hit than others, stresses causing the extinction must be "seeing" traits shared by the species in those groups. Because

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> members of the groups are, by definition, related to each other by common ancestry, they inevitably share some characteristics. For example, all members may have about the same metabolic rate, or size, habitat preference, or geographic range, to mention but a few possibilities. (Raup 1991, 103)

Although Darwin's own narrative model for how to eliminate genera invokes a similar pattern of reasoning, notice that his competition-based scenario cannot explain the simultaneous extinction of many distantly related groups. That is, we can well enough understand the extinction of some genus via its encounter with a wide-ranging, better adapted species or genus. But we cannot explain the simultaneous extinction of several genera, families, orders, and so on, in this way because it would require an extraordinary number of simultaneous large-scale invasions of the right sort. Such an event is exceedingly improbable, and surely a last resort when there are other, much more plausible mechanisms to which we might appeal.

This leaves wanton extinction. If wanton extinction is the only mechanism known to be capable of eliminating lots of higher-order taxa, and if mass extinctions are characterized by the extermination of lots of higher-order taxa, then wanton extinction becomes the best model for understanding how extinction works during these events (it's also a good ecological explanation for the biases in extinction rates). And if wanton extinction explains how mass extinctions work, then wanton extinction also becomes a central part of the story of how major adaptive radiations occur: only wanton extinction eliminates many higher-order taxa, and only the elimination of many higher-order taxa can trigger the aggressive equilibrium response in origination rates that characterizes post–mass extinction "recovery" periods.

But Raup was not content to restrict wanton extinction's purview only to mass extinction events. A number of additional, independent considerations motivated him to argue that this sort of extinction was indeed the most plausible explanation for the general fact that species go extinct:

> Some species die out because they cannot cope in their normal habitat or because superior competitors or predators push them out [i.e., the classical Darwinian model]. But . . . I feel that most species die out because they are unlucky. They die because they are subjected to biological or physical stresses not anticipated in their prior evolution and because time is not available for Darwinian natural selection to help them adapt [i.e., wanton extinction].
>
> Having just made an advocacy statement—bad luck, not bad genes!—I hope the reader appreciates its uncertainties. Favoring bad luck over bad genes is my

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

> best guess. It is shared by many of my colleagues even though a majority of paleontologists and biologists still subscribe to the more Darwinian view of extinction, that of a constructive force favoring the most fit species. (Raup 1991, 191–192)

This passage is significant for our purposes, for two reasons. The first is that wanton extinction is taken here to be a sufficient explanation for extinction in general (with some unusual cases falling under the classical Darwinian model). So the Raup et al. focus on extinction should be understood as more than an attempt to understand some highly significant but also highly unusual events in the history of life (i.e., the five big mass extinctions). The intent is to investigate and potentially completely revise our understanding of why extinctions have ever happened.

The second reason for this passage's significance lies in Raup's open admission of uncertainty with respect to the accuracy of the alternative models. The important implication here is that the nature of extinction is an empirical question, to be decided by the plausibility of each of the models, combined with the ability of each to make sense of the relevant evidence. Once our ability to understand extinction is framed in these terms, the long-lived hegemony of the classical Darwinian model appears to be without justification. Showing that the death of species and higher taxa is an important scientific problem and is rife with uncertainty helped to support paleontology's pretensions toward disciplinary autonomy. I conclude this chapter with an explanation of what exactly it is about death that's so wonderful for paleontology.

9.5 Wonderful Death: The Foundation of Disciplinary Autonomy

Our model of disciplinary autonomy says that the ability of extinction studies to explain the emergence of paleontology as an autonomous theoretical discipline will rest on (1) whether the extinction and phenomena related to extinction were important scientific problems whose solution was in high demand and (2) whether paleontology had a unique ability to disclose the facts of extinction and related phenomena. How does our model perform, measured by these standards?

To the first question, then: How do extinction and related issues rank as generally important scientific problems? We can start to get a sense of their importance by jumping off from the platform provided by Maynard Smith's 1984 reflections in *Nature*, part of which were quoted at the beginning of

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

this chapter. Maynard Smith—as competent a judge as anyone for assessing the importance of a domain of phenomena to our general understanding of the evolutionary process—described paleontology's work on extinction as "the greatest impact that paleontology is having on the way we see the mechanisms of evolution" (Maynard Smith 1984, 402). If he is right, then we should be able to connect the post-MBL work on extinction described in sections 9.3 and 9.4 to certain high-interest features of the evolutionary process in general, and—if we can—to major features of our own evolutionary history.

Recall that Darwin did not find extinction to be a particularly interesting phenomenon in and of itself. Indeed, he was rather nonplussed by the fact that anyone should have ever regarded the death of a species as a puzzling phenomenon. To him it was just obvious that species should go extinct as part of the selective process, and he thus assumed that each case of extinction would be in principle explicable in terms of "conditions . . . becoming less and less favourable," leading at first to rarity and then inevitably to extinction. The causes of extinction were not an open question for Darwin, and he regards his own model as so intuitive (which indeed it is) that he wastes no further time exploring the plausibility of other candidate explanations.

But even though Darwin may not have afforded any intrinsic importance to extinction, it nevertheless plays a number of prominent roles in his thought. Early on it was used to maintain the numerical constancy of species in the face of a tendency in his process of adaptive transmutation to exponentially increase the number of species. Later it was used to explain the fact that organisms are naturally classified as groups within groups. By the time we get to the *Origin*, Darwin applied his favored (read "exclusive") mode of extinction—competitive extermination—to make sense of (1) the absence of intermediate forms both in the fossil record and alive today, (2) "the Forms of Life changing almost simultaneously throughout the World," and (3) the fact that all living beings tend to be near perfection. Thus, even though he could scarcely deign to address the nature of extinction itself, Darwin still elevated extinction to a rank of particular importance by connecting it to a number of phenomena at the center of efforts to understand natural history. Extinction's connection to other high-interest phenomena makes it important.

Likewise, in the current era, extinction now derives much of its importance from its ability to explain historically significant taxonomic radiations

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

and replacements—a connection of which early extinction and diversity theorists were well aware and that Maynard Smith (1984, 402) cites as largely responsible for paleontology's impact "on the way we see the mechanisms of evolution." The sheer volume of organismic diversity—standing diversity as well as diversity across time—is a rather fascinating feature of our world and a particularly surprising feature of the living world if one approaches that world from the perspective of Darwinian theory. Recall what Raup had written to Schopf in the letter quoted above: If Darwin is right about extinction and diversity, we ought to not see the fossil patterns that we do—specifically, the steady growth in diversity over 540 million years. Indeed, many of the principal participants in the mathematical development of Darwinian theory during the 1920s–1950s concluded that natural populations should exhibit little to no diversity (Wallace 1991; Gayon 1998). It therefore came as quite a shock when the first efforts to measure genetic diversity in the wild revealed an immense amount of natural variation (Lewontin and Hubby 1966; see Wallace 1991 and Gayon 1998 for discussion). By providing a simple, general, and intuitive model for how extinction rates might affect diversification patterns over millions of years, paleontologists were able to offer a higher-order way of organizing critically important but frustratingly truculent biological facts. Extinction thus becomes important because it is the phenomenon that mediates our ability to understand part of the central and difficult problem of organismic diversity.

Extinction's degree of scientific import has also benefited from the fundamental revisions that have occurred surrounding our understanding of taxon death per se, an advantage that it has not enjoyed since the publication of the *Origin*. The revelation that extinction itself probably does not—cannot—work in the way it had been assumed to work for the past 125 years was, as Maynard Smith found it, "surprising," and delivered extinction from the dusty crate of well-understood humdrum to the iridescent pile of phenomena that cry out for explanation. Paleontologists' efforts to reopen the book on the nature of extinction have given it an intrinsic importance that it holds quite independently of anything extinction can be used to explain, in the same way that we would regard as worthy of our attention any other significant feature of our world that we suddenly discovered we did not understand. Extinction is a central part of the evolutionary process. If we don't understand extinction, then there is a lot we don't understand about evolution.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

This is where the satisfaction of the Darwinian population's *continuity* and *heredity* criteria is most easily discerned. The 1970s witnessed a major shift in the research focus of paleontologists toward an emphasis on problems associated with extinction, and a major shift in the approach to these problems (exemplified by the species-as-particles framework). These major shifts were, I have argued, connected with an equally significant change in the general recognition of extinction as a scientific problem that paleontologists were uniquely well suited to explore. This transition was initiated by the early explorations of the MBL group and a handful of like-minded thinkers, such as James Valentine and Steven Stanley. By the late 1970s–early 1980s, paleontologists inspired and trained by these researchers were using the modestly refined versions of species-as-particles framework to make commensurately modest degrees of progress on slightly refined versions of the initial problems explored in MBL. Thus they were confirming Raup's 1979 prediction that "scores of young paleontologists will be plotting survivorship curves (or whatever)" in the event that he, Schopf, and Gould "really succeed in selling our current brand of nomothetic paleontology."[1] This corresponds well to Orr's analysis of the adaptive walk, in which early changes tend to be larger and account for a sizable chunk of the overall fitness gain. Subsequent changes are minor refinements on the direction set by earlier moves and result in comparatively minor fitness gains.

With extinction we can also see the increasing degree to which the *intrinsicality* condition came to be satisfied. Because paleontology was in a unique position to disclose and explain the facts of extinction and important related phenomena, paleontology in general, and species-as-particles in particular, offers us a higher-order way of organizing critically important but frustratingly truculent biological facts that could not even be addressed at the population genetic level, let alone understood. In this way, the success of paleontological practices would come more and more to depend on their ability to contribute to research problems that were the sole property of paleontology. Consider, for example, that obvious fact that most of the empirical details that can inform our understanding of extinction are deposited in the fossil record. This basic biological truism already puts paleontologists in a privileged position vis-à-vis extinction: only paleontologists have direct access to the record of actual species deaths. They are thus uniquely situated to disclose the facts concerning what has died, as well as where and when. This unique position consequently privileges

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

paleontologists when it comes to the business of formulating and testing theories about extinction, be it extinction as a general evolutionary mechanism or the extinction of particular taxa. Any sensible theory of extinction will have to make sense of the fossil record. As we have mentioned previously, Darwin's theory of extinction is supremely intuitive, thoroughly general, and fundamentally related to an even more general picture of evolution—that is, natural selection. But as Raup pointed out early on, Darwin's theory of extinction cannot make sense of the fossil record. Given that the fossil record is the authoritative repository of facts about extinction, this counts as a major strike against Darwin's view.

This bit about facts and theory testing should be uncontroversial. Is it equally obvious that paleontologists' privileged perspective on the fossil record underwrites a similarly privileged ability to formulate theories of extinction—in other words, to develop strategies for solving problems associated with extinction? This is an important question, particularly in light of what Gould and many others saw as the Master Argument for paleontology's autonomy—the reclamation of paleontologists' "proper role . . . as generators of testable and irreducible theories about pattern in the history of life." Why should paleontologists' ability to corner the fossil market translate into a unique ability to generate "testable and irreducible theories about pattern in the history of life"? Can't anyone do this, whether she understands fossils or not?

In principle, yes. For example, a population geneticist could develop a theory that describes the dynamics of mass extinction at the genetic level. Were such a theory successful, it could properly be said to have reduced the macroevolutionary dynamics of mass extinction to the microevolutionary dynamics of allelic replacement. But there is no such theory, successful or unsuccessful. Nor is there likely to be. Population genetics is designed to explain short-term changes in populations via reference to differences in the fitnesses of different alleles. This framework would make population genetics well suited to explain extinction if extinction were predominantly a matter of the steady competitive replacement of some taxa by other taxa that are better adapted to the local conditions. But extinction does not appear to work this way. Rather, if Raup and company are anywhere near correct, extinction works primarily through species' inability to cope with sudden radical changes in the physical environment. A species' relationship to its physical environment is not something population genetics is

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

in the business of explaining. It is only equipped for telling us what will happen in a population, given that it is composed of such-and-such types of organisms that bear such-and-such relationships to their environments (i.e., their fitness). To be sure, that is a very useful thing to be able to do. But it does not make population genetics particularly fit for telling us why the fossil record exhibits the patterns of extinction and diversity that we see.

Does this mean that macroevolutionary dynamics are not reducible to microevolutionary dynamics describable by population genetics? No. I suppose all it means is that we currently cannot see how to account for the large-scale facts of extinction and diversity by appealing to allelic fitness differences. That is not sufficient for irreducibility, which is a claim about the necessary nonexistence of a particular mapping function. I don't think we have the evidence required to establish that sort of thing, even in the case of something like Jablonski's Law, which appears to not supervene on lower-level species facts. But while it may not be sufficient for irreducibility, it is also not necessary for autonomy. Whether or not there is a particular mapping function out there waiting to be discovered by some genius is really beside the point. What matters for autonomy is the fact that the species-as-particles approach is the only game in town. Paleontology's perspective on extinction and diversity is uniquely well suited for solving those problems, not because it is the only possible perspective on these problems, but because it is, in fact, the only perspective on these problems. Whether its perspective is necessarily the only possible one, or instead by sheer historical accident happens to be the only one currently available, it is unique in any case. The details of why it happens to be the unique perspective are simply irrelevant to whether it is unique. For my part, I do think that paleontology's privilege is well protected, in the sense that a reduction of the macroevolutionary theories of extinction and diversity to allelic fitness differences strikes me as beyond improbable. Of course, this does not show that paleontology's vise-grip on the explanation of these phenomena is necessary because the theories are irreducible; it is merely an expression of the suspicion that paleontology's monopoly is likely to persist. But whether that is true is irrelevant to whether there is a legitimate basis for disciplinary autonomy in this case. What is relevant is that paleontology is in possession of a theoretical and evidential base that, for whatever reason, is the only available means for making sense of a particular domain of phenomena.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

It is a particular strength of our model that it does not need to appeal to the concept of irreducibility in any way in order to explain perceptions of legitimate claims to disciplinary autonomy. An analogy with consumer goods markets makes this clear. When some firm holds a monopoly over the production of certain goods, the details of how that monopoly is maintained do not help us to understand the firm's ability to control prices. The firm's ability to control prices is explained simply by the fact that, for whatever reason, the firm has a monopoly. The effects of a monopoly are not dependent on further facts about that monopoly, nor are they dependent on facts concerning the possibility of other firms—someday, somehow—invading the market. Similarly, the legitimacy of paleontology's autonomy is explained by the mere fact that paleontology monopolized the explanatory resources for understanding extinction and diversity. In these domains, paleontology could do what no other group of researchers could do: disclose, predict, and explain the facts related to patterns in extinction and diversity across time.

9.6 Conclusion

We have looked at four episodes in the recent history of paleontology with an eye toward understanding its development as an autonomous scientific discipline. Using the evolutionary model of disciplinary emergence, I argued that we ought to be able to explain paleontology's disciplinary autonomy in terms of (1) the identification of a set of scientific problems whose solution was in high demand and (2) the availability of a privileged perspective that is uniquely well suited for solving those problems.

The first episode looked at Stephen Jay Gould's attempt to forge a "science of form" over which paleontology would exercise dominion. This did not end well for Gould, and the evolutionary model was able to explain why: quite simply, the good design of organisms was already well understood. But some good did come out of Gould's efforts—namely, a more demanding standard for demonstrating adaptation and an interest in causes of suboptimal design.

The second episode looked broadly at the notion of punctuated equilibrium, which Gould defended for most of his career as a major contribution to our general understanding of the evolutionary process and thus part of the foundation for paleontology's disciplinary autonomy. Our model

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

suggests, correctly I believe, that this cannot be true. For one, the phrase "punctuated equilibrium" refers to (at best) a pseudoprocess and thus fails to identify a genuine scientific problem. In addition, the appearance of this pseudoprocess is explicable in population genetic terms. So, to the extent that there even was a distinctively paleontological proposal for how to understand the appearance of this pseudoprocess, it would not have been a privileged one.

The third episode followed the rise of the species-as-particles problem-solving strategy, looking in detail at its development and application in the work of the MBL group. The MBL group's contribution, although central to the eventual emergence of paleontology as an autonomous discipline, could not in and of itself serve as the basis for disciplinary autonomy. From the evolutionary model's perspective, this was because the MBL group failed to disclose any real scientific problem. The principal phenomena with which the MBL group was working were simulated clades, and even for simulated clades, they were unrealistic. The major payoff for the MBL effort was the development of an entirely unique approach to interpreting the fossil record and to studying evolutionary processes more generally: the species-as-particles approach.

The species-as-particles problem-solving strategy would eventually allow paleontologists to disclose and explain phenomena of critical evolutionary importance, phenomena that are neither predicted nor explained by other strategies for solving evolutionary problems. Many of the large-scale puzzles related to extinction and diversity are solved by assuming that the peculiarities of individual taxa do not matter. No other explanatory framework has achieved measurable success in these domains, which the evolutionary model explains by reference to paleontology's unique theoretical and evidential perspective.

Another goal of this study was to see how we might resolve Gould's "paradox": on the one hand, he believed that all biological facts would have been otherwise were we to rewind the tape of life; on the other hand, he believed that there were "nomothetic undertones" to the evolutionary process, which suggests that at least some biological facts could not have been otherwise—a contradiction! Gould was a creature of conflict, and this description of two of the dominant threads in his thought does not ease the tension in any way. But I think the problem is simply that the description is too coarse-grained. Gould believed that the particular facts about

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

what forms happen to have appeared on Earth would have been otherwise, were we to replay the tape of life. This was a view he held at least since the publication of his "Dollo on Dollo's Law" paper (Gould 1970a), where he describes the species level as so fraught with evolutionary noise as to be completely unhelpful in understanding general properties of evolutionary systems. But he did not believe that this noise necessarily percolated up to higher taxonomic levels. For example, many of the facts about clades, he believed, would have held no matter what. In keeping with Gould's metaphor, even if replaying the tape of life produces radically different biological forms, the tape always winds through the machine in the same way. We saw Raup's work on extinction and Sepkoski's work on diversity as particularly mature expositions of this view. These research programs each explained important components of the evolutionary process in terms of mechanisms that hold independently of the lower-level biological details. Thus, what looks like miscellany at one level looks highly structured, even necessary, at another level. This "fruitful lawmaking strategy" originally sought—and in no small part developed—by Gould has in many ways come to define the field on whose behalf he so tirelessly campaigned.

Looking at the history as I have presented it here (and in much more detail as presented in D. Sepkoski 2012), I think it is correct to say that of the two ingredients required to catalyze the emergence of the new discipline of evolutionary paleontology, the species-as-particles problem-solving approach was by far the more significant; extinction and diversity, although unquestionably of enduring evolutionary importance, were more or less in the right place at the right time. This is as it should be. Independently of the considerations brought out by the evolutionary model, we would expect the defining feature of an autonomous discipline to be its "lawmaking strategy" or the general structure of its theories, rather than the particular problem set with which it happens to be grappling at any given time.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Notes

Introduction

1. This is the subtitle of Shapin's book (2010).

2. See Hull (1988, 320), where he explains how scientists adhere to the community norms through the combined pressures of reward for doing so and punishment for not doing so.

Chapter 1

1. He is still arguing for it as late as the monumental "World Changes" conference (Kuhn 1993).

2. It is important to add here that the value of Popper's and Campbell's evolutionary themata does not stand or fall with how productive their use of the analogy is. Campbell's studies in particular are extremely rich.

3. Boaz Miller (2012) makes a related point about scientific consensus.

4. Is it possible that sociocultural and epistemic pressures will overlap—that is, is it possible that there will be cases where an improvement to our knowledge is causally explained by attempts to meet a sociocultural goal? Of course, this is possible. And were these categorically different sorts of pressures to overlap completely, such that they were indistinguishable from one another, it would be difficult to argue persuasively for an epistemic interpretation of the directionality of science (I don't think they are indistinguishable).

5. Godfrey-Smith (2009, 53). "Intrinsicality" is my word, not Godfrey-Smith's.

6. My thanks to an anonymous referee for encouraging me to emphasize this important connection.

7. This concern was raised by an anonymous referee.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

8. Toulmin's account, for example, included an evolutionary tree for a hypothetical concept that was subject to different kinds of evolutionary changes (e.g., branching, extinction) depending on environmental factors. See Toulmin (1972, 200, 203, 205).

Chapter 2

1. Haufe (2013) refines Brandon's discussion by articulating the assumption of necessity inherent in fitness assignments.

2. The pressure to suppress variation is discussed in chapter 3.

3. Miller (2012) describes different causal processes that might lead to scientific consensus. These can be instructively treated as different selection regimes.

4. "Why not have arranged marriages, too?" replied a historian to whom I described this custom.

5. See, for example, the list of permissible particle physics thesis topics for a DPhil in physics at Oxford: https://www.physics.ox.ac.uk/study/postgraduates/dphil-particle-physics/running-experiments.

6. Cf. Sterelny (2012). Usselman et al. (2005) describe the extraordinary volume of "gestural knowledge" that escaped Justus Liebig's published description of his Kaliapparat (see section 2.3.1). The locus classicus for the study of tacit knowledge is Polanyi (1958).

7. A good description is contained in De Gandt (1995, chap. 1).

8. See Boyd et al. (2011, 10922) for an exploration of this form of pedagogy as a cultural evolutionary phenomenon.

9. Holmes cites the French chemist Jean-Baptiste Dumas's opinion that the discrepancies were due not to imprecise measurement but to impure samples.

10. Berzelius, *Jahres-Bericht über die Fortschritte der physischen Wissenschaften*, vol. 13 (Tübingen: Laupp, 1834), 186–188, quoted in Rocke (1984, 174).

11. Rocke 2003, 94. These differences arise because of the promiscuity enabled by covalent bonding, which was unknown at the time.

Chapter 3

1. Time to optimization in turn depends on things like rate of gene flow, mutation rate, size of population, and strength of selection—all of which determine how quickly the population sifts through its various morphological options.

2. Beginning with Lande (1981).

3. "RFA-AI-17–001," https://grants.nih.gov/grants/guide/rfa-files/RFA-AI-17-001.html, accessed February 9, 2017.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

4. http://www.niaid.nih.gov/researchfunding/grant/strategy/pages/3researchstrategy.aspx#f01, accessed August 3, 2012.

5. See Polanyi (1958) and Kuhn (1977) for an extended discussion and analysis.

6. See, e.g., Simonton (1979); Bettencourt et al. (2006); and Bettencourt, Kaiser, and Kaur (2009).

7. Thanks to Kevin Laland for helpful discussion on this point.

8. This observation has previously been pointed out by Hull (1988), Kitcher (1993), and Goldman (1999), among others.

Chapter 4

1. The connection between this passage and cultural evolution is explored in Mesoudi, Whiten, and Laland (2004).

2. Shennan (2013, 152) provides a survey of recent attempts to model the evolution of technology using fitness landscapes.

3. Orr's simulations suggested an initial step of about 30% of total fitness gain made during an adaptive walk.

4. The limit case would be where there is complete overlap, and the optima are superimposed. A movement of arbitrary magnitude toward one optimum necessarily travels the same distance toward all.

5. It is worth emphasizing that, in contrast both to the contributions that precede and that follow them, most of the epochal, paradigm-inducing achievements of science are not distinguished by their contributions to our store of empirical knowledge (*Opticks* is clearly an exception). Cohen (1983, 149), for example, argued forcefully that the Newtonian Revolution wrought by the *Principia* consisted of Newton's having "set forth a style of science that shows how mathematical principles might be applied to physics and astronomy (that is, to natural philosophy) in a particularly fruitful way." Darwin's *Origin* is similar in that it does not report any original discoveries (other than some very low-level results concerning things like seed dispersal and abundances of various types of plants). More than anything, what made the book such a monumental contribution was the way in which it organized existing knowledge and demonstrated an effective framework for future inquiry (Kitcher 1985).

6. From the documentary movie *Particle Fever* (2014, directed by Mark Levinson).

Chapter 5

1. Darwin (1859) engages in this latter exercise in several places in the *Origin*. See, for example, his discussion of "reversion" in chapter 2.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

2. I use the broad term "paleontology" to refer to a variety of subdisciplines in evolutionary biology that study the history of life on relatively large time scales.

3. Namely, (8). But (8) is nonstandard. It is true that the term "successful" is often attached to lineages that survive a mass extinction, but it would be unusual to think of surviving a mass extinction as single-handedly constitutive of evolutionary success. One could of course (as practitioners often do) say that a lineage "successfully survived a mass extinction." But such a lineage would not thereby be evolutionary successful.

4. Thank you to Carl Simpson for helpful discussion on this point.

5. Raup (1991). It is worth mentioning that Van Valen (1973) derived the surprising result that longevity does not reduce the probability of extinction.

6. Darwin (1859). Regrettably, Darwin's deep concern with the "largeness" of species and genera has not been given the scholarly attention it deserves. Browne (1980) and Ospovat (1983) are important exceptions.

7. Although, of course, it can be causally related. Atkinson et al. (2008) investigate the emergence of languages as an evolutionary branching process.

8. Bürger (2005).

9. Gavrilets (2005) criticizes mate preference models of ecological speciation for building in strong assumptions that tend to provide little understanding of ecological speciation per se. For what it's worth, I agree with his position.

Chapter 6

1. Raup and Sepkoski (1984) proposed that mass extinction events on Earth may be the result of a star with a 27 Myr orbit. The 27 Myr period has been rejected, but the periodicity of mass extinctions itself remains a live hypothesis. See Melott and Bambach (2014).

2. J. Sepkoski (1984) showed this cycle to be a property of any system characterized by an equal probability of (1) reproducing, (2) branching, (3) going extinct, and (4) being subject to random high-magnitude perturbations.

3. Newell (1963, 1967).

4. Simpson (1944). This characterization comes from Erwin and Valentine (2013, 228).

5. Many thanks to John Huss for referring me to this paper.

6. Godfrey Smith (2009, 60–62) undertakes such a feat in his discussion of drift.

7. Kuhn (1962b, 155–156). For further thoughts on these kinds of decisions, see Kuhn (1977, chap. 13).

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

8. I explore the many dimensions of fruitfulness in Haufe (2023). The general evolutionary significance of fruitfulness will be the focus of future work.

9. See Brandon (2005) and Godfrey-Smith (2009, 59–63) for more on the relative significance of different evolutionary factors under different conditions.

10. See, for example, Kuhn's "Postscript-1969" (1970, 206).

11. See Wray (2011) for a helpful discussion of Kuhn's vexed list of revolutions.

12. This process of lexical refinement is central to the notion of progress that Kitcher (1993) defends.

13. I employ Wray's (2011) use of this term (at least, I think I do).

14. Personal communication, "Chance in Evolution" conference, Duke University, April 2007.

15. See, e.g., Tappenden (2005); Grosholz (2007).

Chapter 7

1. John Beatty (personal communication) has described the sentiments expressed by Gould in "Is Uniformitarianism Necessary?" (1965) as being reflective of some of the contents of Gould's papers from his undergraduate days at Antioch.

Chapter 9

1. Raup to Schopf, January 28, 1979, Schopf pap. 3, 30, Smithsonian Institution Archives, Thomas J. M. Schopf Papers, Washington DC, quoted in D. Sepkoski (2012), 259.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

References

Abbott, A. 2016. Structure as Cited, Structure as Read. In *Kuhn's* Structure of Scientific Revolutions at Fifty: *Reflections on a Science Classic*, edited by R. Richards and L. Daston, 167–181. Chicago: University of Chicago Press.

Atkinson, Q. D., A. Meade, C. Venditti, S. J. Greenhill, and M. Pagel. 2008. Languages Evolve in Punctuational Bursts. *Science* 319(5863): 588.

Bardeen, C. G., R. R. Garcia, O. B. Toon, and A. J. Conley. 2017. On Transient Climate Change at the Cretaceous–Paleogene Boundary Due to Atmospheric Soot Injections. *Proceedings of the National Academy of Sciences of the United States of America* 114(36): E7415–E7424.

Beatty, J. 1995. The Evolutionary Contingency Thesis. In *Concepts, Theories, and Rationality in the Biological Sciences*, edited by G. Wolters and J. G. Lennox, 45–81. Pittsburgh: University of Pittsburgh Press.

Beatty, J. 2006. Replaying Life's Tape. *Journal of Philosophy* 103: 336–362.

Bettencourt, L. M. A., A. Cintrón-Arias, D. I. Kaiser, and C. Castillo-Chávez. 2006. The Power of a Good Idea: Quantitative Modeling of the Spread of Ideas from Epidemiological Models. *Physica A: Statistical Mechanics and Its Applications* 364: 513–536.

Bettencourt, L. M. A., D. I. Kaiser, and J. Kaur. 2009. Scientific Discovery and Topological Transitions in Collaboration Networks. *Journal of Informetrics* 3(3): 210–221.

Block, N. 1986. Advertisement for a Semantics for Psychology. *Midwest Studies in Philosophy* 10(1): 615–678.

Bloor, D. 1974. Popper's Mystification of Objective Knowledge. *Science Studies* 4(1): 65–76.

Bolnick, D. I., and B. M. Fitzpatrick. 2007. Sympatric Speciation: Models and Empirical Evidence. *Annual Review of Ecology, Evolution, and Systematics* 38(1): 459–487.

Bonança, M. V. S., and C. Jarzynski. 2016. Conditional Reversibility in Nonequilibrium Stochastic Systems. *Physical Review E* 93(2): 022101.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Boyd, R., P. J. Richerson and J. Henrich. 2011. The Cultural Niche: Why Social Learning Is Essential for Human Adaptation. *Proceedings of the National Academy of Sciences of the United States of America* 108: 10918–10925.

Brandon, R. 1978. Adaptation and Evolutionary Theory. *Studies in the History and Philosophy of Science* 9(3): 181–206.

Brandon, R. N. 2005. The Difference between Selection and Drift: A Reply to Millstein. *Biology and Philosophy* 20(1): 153–170.

Browne, J. 1980. Darwin's Botanical Arithmetic and the "Principle of Divergence," 1854–1858. *Journal of the History of Biology* 13(1): 53–89.

Brush, S. G. 2009. *Choosing Selection: The Revival of Natural Selection in Anglo-American Evolutionary Biology, 1930–1970*. Philadelphia: American Philosophical Society.

Bürger, R. 2005. A Multilocus Analysis of Intraspecific Competition and Stabilizing Selection on a Quantitative Trait. *Journal of Mathematical Biology* 50(4): 355–396.

Burkhardt, R. 1977. *The Spirit of System: Lamarck and Evolutionary Biology*. Cambridge, MA: Harvard University Press.

Camp, E. 2019. Imaginative Frames for Scientific Inquiry: Metaphors, Telling Facts, and Just-So Stories. In *The Scientific Imagination*, edited by A. Levy and P. Godfrey-Smith, 304–336. New York: Oxford University Press.

Campbell, D. T. 1974. Evolutionary Epistemology. In *The Philosophy of Karl Popper*, edited by P. Schlipp, 413–463. La Salle, IL: Open Court.

Campbell, D. T. 1977. Comment on "The Natural Selection Model of Conceptual Evolution." *Philosophy of Science* 44(3): 502–507.

Carnap, R. 1950. Empiricism, Semantics, and Ontology. *Revue Internationale de Philosophie* 4: 20–40.

Cloud, P. E., Jr. 1948. Some Problems and Patterns of Evolution Exemplified by Fossil Invertebrates. *Evolution* 2(4): 322–350.

Cohen, I. B. 1983. *The Newtonian Revolution*. New York: Cambridge University Press.

Collins, H. 1992. *Changing Order: Replication and Induction in Scientific Practice*. Chicago: University of Chicago Press.

Darwin, C. R. 1835. The Position of the Bones of Mastodon (?) at Port St Julian Is of Interest. CUL-DAR42.97-99. Transcribed and edited by John van Wyhe. Darwin Online, http://darwin-online.org.uk/.

Darwin, C. R. 1837–1838. Notebook B: [Transmutation of Species]. CUL-DAR121. Transcribed by Kees Rookmaaker. Darwin Online, http://darwin-online.org.uk/.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Darwin, C. R. 1838. Notebook C: [Transmutation of Species]. CUL-DAR122. Transcribed by Kees Rookmaaker. Darwin Online, http://darwin-online.org.uk/.

Darwin, C. R. 1845. Forbes Widest Distributed Forms Thinks Probably Oldest Created? CUL-DAR205.4.40-42. Edited by John van Wyhe. Darwin Online, http://darwin-online.org.uk/.

Darwin, C. 1859. *On the Origin of Species*. London: John Murray.

Dennett, D. C. 1995. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York: Simon & Schuster.

Dobzhansky, T. 1937. *Genetics and the Origin of Species*. New York: Columbia University Press.

Donoghue, M. J. 2005. Key Innovations, Convergence, and Success: Macroevolutionary Lessons from Plant Phylogeny. *Paleobiology* 31(2): 77–93.

Eblen, M. K., R. M. Wagner, D. Roychowdhury, K. C. Patel, and K. Pearson. 2016. How Criterion Scores Predict the Overall Impact Score and Funding Outcomes for National Institutes of Health Peer-Reviewed Applications. *PLOS One* 11(6): e0155060.

Eldredge, N. 1971. The Allopatric Model and Phylogeny in Paleozoic Invertebrates. *Evolution* 25(1): 156–167.

Eldredge, N. 2008. The Early "Evolution" of "Punctuated Equilibria." *Evolution: Education and Outreach* 1(2): 107–113.

Eldredge, N., and S. J. Gould. 1972. Punctuated Equilibria: An Alternative to Phyletic Gradualism. *Models in Paleobiology*: 82–115.

Eldredge, N., J. N. Thompson, P. M. Brakefield, S. Gavrilets, D. Jablonski, J. B. C. Jackson, R. E. Lenski, B. S. Lieberman, M. A. Mcpeek, and W. Miller III. 2005. The Dynamics of Evolutionary Stasis. *Paleobiology* 31(2): 133–145.

Eren, O., and N. Mocan. 2018. Emotional Judges and Unlucky Juveniles. *American Economic Journal: Applied Economics* 10(3): 171–205.

Erwin, D. H. 2006. *Extinction: How Life on Earth Nearly Ended 250 Million Years Ago*. Princeton, NJ: Princeton University Press.

Erwin, D. H., and J. W. Valentine. 2013. *The Cambrian Explosion: The Construction of Animal Biodiversity*. Greenwood Village, CO: Roberts and Company.

Erwin, D. H., J. W. Valentine, and J. Sepkoski. 1987. A Comparative Study of Diversification Events: The Early Paleozoic versus the Mesozoic. *Evolution* 41(6): 1177–1186.

Feder, J. L., S. P. Egan, and P. Nosil. 2012. The Genomics of Speciation-with-Gene-Flow. *Trends in Genetics* 28(7): 342–350.

Fisher, R. 1930. *The Genetical Theory of Natural Selection*. Oxford: Clarendon Press.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Forman, P. 1971. Weimar Culture, Causality, and Quantum Theory, 1918–1927: Adaptation by German Physicists and Mathematicians to a Hostile Intellectual Environment. *Historical Studies in the Physical Sciences* 3: 1–115.

Fry, J. D. 2003. Multilocus Models of Sympatric Speciation: Bush versus Rice versus Felsenstein. *Evolution* 57(8): 1735–1746.

Futuyma, D. 1987. On the Role of Species in Anagenesis. *The American Naturalist* 130(3): 465–473.

Galison, P. 2016. Practice All the Way Down. In *Kuhn's* Structure of Scientific Revolutions at Fifty: *Reflections on a Science Classic*, edited by R. Richards and L. Daston, 42–70. Chicago: University of Chicago Press.

Gandt, F. de. 1995. *Force and Geometry in Newton's Principia*. Princeton, NJ: Princeton University Press.

Gavrilets, S. 2004. *Fitness Landscapes and the Origin of Species*. Princeton, NJ: Princeton University Press.

Gavrilets, S. 2005. "Adaptive Speciation"—It Is Not That Easy: Reply to Doebeli et al. *Evolution* 59(3): 696–699.

Gayon, J. 1998. *Darwinism's Struggle for Survival*. New York: Cambridge University Press.

Gillespie, J. H. 1984. Molecular Evolution over the Mutational Landscape. *Evolution* 38(5): 1116–1129.

Godfrey-Smith, P. 2003. *Theory and Reality*. Chicago: University of Chicago Press.

Godfrey-Smith, P. 2009. *Darwinian Populations and Natural Selection*. New York: Oxford University Press.

Godfrey-Smith, P. 2012. Darwinism and Cultural Change. *Philosophical Transactions of the Royal Society B: Biological Sciences* 367(1599): 2160–2170.

Goldman, A. I. 1999. *Knowledge in a Social World*. New York: Oxford University Press.

Goodman, N. 1968. *Languages of Art: An Approach to a Theory of Symbols*. Indianapolis: Bobbs-Merrill.

Gould, S. J. 1965. Is Uniformitarianism Necessary? *American Journal of Science* 263(3): 223–228.

Gould, S. J. 1966. Allometry and Size in Ontogeny and Phylogeny. *Biological Reviews* 41(4): 587–638.

Gould, S. J. 1970a. Dollo on Dollo's Law: Irreversibility and the Status of Evolutionary Laws. *Journal of the History of Biology* 3(2): 189–212.

Gould, S. J. 1970b. Evolutionary Paleontology and the Science of Form. *Earth-Science Reviews* 6(2): 77–119.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Gould, S. J. 1973. The Shape of Things to Come. *Systematic Zoology* 22(4): 401–404.

Gould, S. J. 1974. The Origin and Function of "Bizarre" Structures: Antler Size and Skull Size in the "Irish Elk," *Megaloceros Giganteus*. *Evolution* 28(2): 191–220.

Gould, S. J. 1978. Sociobiology: The Art of Storytelling. *New Scientist* 80(1129): 530–533.

Gould, S. J. 1980. The Promise of Paleobiology as a Nomothetic, Evolutionary Discipline. *Paleobiology* 6(1): 96–118.

Gould, S. J. 1985. The Paradox of the First Tier: An Agenda for Paleobiology. *Paleobiology* 11(1): 2–12.

Gould, S. J. 1988. On Replacing the Idea of Progress with an Operational Notion of Directionality. In *Evolutionary Progress*, edited by M. H. Nitecki, 319–338. Chicago: University of Chicago Press.

Gould, S. J. 1989. *Wonderful Life: The Burgess Shale and the Nature of History*. New York: W. W. Norton.

Gould, S. J. 2002. *The Structure of Evolutionary Theory*. Cambridge, MA: Belknap Press of Harvard University Press.

Gould, S. J. 2007. *Punctuated Equilibrium*. Cambridge, MA: Belknap Press of Harvard University Press.

Gould, S. J., and N. Eldredge. 1977. Punctuated Equilibria: The Tempo and Mode of Evolution Reconsidered. *Paleobiology* 3(2): 115–151.

Gould, S. J., and R. Lewontin. 1979. The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme. *Proceedings of the Royal Society of London—Series B: Biological Sciences* 205(1161): 581–598.

Gould, S. J., D. M. Raup, J. Sepkoski Jr., and T. Schopf. 1977. The Shape of Evolution: A Comparison of Real and Random Clades. *Paleobiology* 3(1): 23–40.

Grosholz, E. 2007. *Representation and Productive Ambiguity in Mathematics and the Sciences*. New York: Oxford University Press.

Gutting, G. 2009. *What Philosophers Know: Case Studies in Recent Analytic Philosophy*. New York: Cambridge University Press.

Hacking, I. 1979. Imre Lakatos's Philosophy of Science. *British Journal for the Philosophy of Science* 30(4): 381–402.

Hacking, I. 1993. Working in a New World: The Taxonomic Solution. In *World Changes: Thomas Kuhn and the Nature of Science*, edited by P. Horwich, 275–310. Cambridge, MA: MIT Press.

Hacking, I. 2012. Introductory Essay. In T. Kuhn, *The Structure of Scientific Revolutions*, 4th ed., vii–xxxvii. Chicago: University of Chicago Press.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Haufe, C. 2013. From Necessary Chances to Biological Laws. *British Journal for the Philosophy of Science* 64(2): 279–295.

Haufe, C. 2015. Gould's Laws. *Philosophy of Science* 82(1): 1–20.

Haufe, C. 2023. *Fruitfulness*. New York: Oxford University Press.

Haufe, C. ms-a. The Evolutionary Fixation of Ideas.

Haufe, C. ms-b. Funding Fruitfulness.

Heideman, M. T., D. H. Johnson, and C. S. Burrus. 1985. Gauss and the History of the Fast Fourier Transform. *Archive for History of Exact Sciences* 34(3): 265–277.

Holmes, F. L. 1989. The Complementarity of Teaching and Research in Liebig's Laboratory. *Osiris* 5: 121–164.

Holton, G. 1962. Models for Understanding the Growth and Excellence of Scientific Research. In *Excellence and Leadership in a Democracy*, edited by S. R. Graubard and G. Holton, 94–131. New York: Columbia University Press.

Hoyningen-Huene, P. 1993. *Reconstructing Scientific Revolutions: Thomas S. Kuhn's Philosophy of Science*. Chicago: University of Chicago Press.

Hull, D. 1988. *Science as a Process*. Chicago: University of Chicago Press.

Huss, J. 2009. The Shape of Evolution: The MBL Model and Clade Shape. In *The Paleobiological Revolution*, edited by D. Sepkoski and M. Ruse, 327–345. Chicago: University of Chicago Press.

Jablonski, D. 1986. Background and Mass Extinctions: The Alternation of Macroevolutionary Regimes. *Science* 231(4734): 129–133.

Jablonski, D. 2002. Survival without Recovery after Mass Extinctions. *Proceedings of the National Academy of Sciences of the United States of America* 99(12): 8139–8144.

Jablonski, D. 2005. Mass Extinctions and Macroevolution. *Paleobiology* 31(2): 192–210.

Janiak, A. 2007. Newton and the Reality of Force. *Journal of the History of Philosophy* 45(1): 127–147.

Kimura, M. 1983. *The Neutral Theory of Molecular Evolution*. New York: Cambridge University Press.

Kitcher, P. 1985. Darwin's Achievement. In *Reason and Rationality in Natural Science*, edited by N. Rescher, 127–189. Lanham, MD: University Press of America.

Kitcher, P. 1993. *The Advancement of Science*. New York: Oxford University Press.

Kitcher, P. 2015. Pragmatism and Progress. *Transactions of the Charles S. Peirce Society* 51(4): 475–494.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Knight, J. B. 1947. Paleontologist or Zoologist. *Bulletin of the Geological Society of America* 58(4): 281–286.

Kohn, D. 1980. Theories to Work By: Rejected Theories, Reproduction, and Darwin's Path to Natural Selection. *Studies in the History of Biology* 4: 67–170.

Kohn, D. 2008. Darwin's Keystone: The Principle of Divergence. In *The Cambridge Companion to the Origin of Species*, edited by M. Ruse and J. R. Richards, 87–108. New York: Cambridge University Press.

Kuhn, T. S. 1961. The Function of Measurement in Modern Physical Science. *Isis* 52(2): 161–193.

Kuhn, T. S. 1962a. Historical Structure of Scientific Discovery. *Science* 136(3518): 760–764.

Kuhn, T. S. 1962b. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Kuhn, T. S. 1970. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Kuhn, T. S. 1977. Second Thoughts on Paradigms. In *The Essential Tension*, 293–319. Chicago: University of Chicago Press.

Kuhn, T. S. 1982. Commensurability, Comparability, Communicability. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1982: 669–688.

Kuhn, T. S. 1993. Afterwords. In *World Changes: Thomas Kuhn and the Nature of Science*, edited by P. Horwich, 311–341. Cambridge, MA: MIT Press.

Kuhn, T. S. 2000. *The Road since* Structure: *Philosophical Essays, 1970–1993, with an Autobiographical Interview*. Edited by J. Conant and J. Haugeland. Chicago: University of Chicago Press.

Laland, K. N., J. Odling-Smee, and M. W. Feldman. 2000. Niche Construction, Biological Evolution, and Cultural Change. *Behavioral and Brain Sciences* 23(1): 131–146.

Lande, R. 1981. Models of Speciation by Sexual Selection on Poly Genic Traits. *Proceedings of the National Academy of Sciences of the United States of America* 78(6): 3721–3725.

Laudan, L. 1977. *Progress and Its Problems*. Berkeley: University of California Press.

Laudan, R. 1993. Histories of the Sciences and Their Uses: A Review to 1913. *History of Science* 31(91): 1–34.

Lewontin, R. 1985. Population Genetics. *Annual Review of Genetics* 19(1): 81–102.

Lewontin, R. C., and J. L. Hubby. 1966. A Molecular Approach to the Study of Genic Heterozygosity in Natural Populations. II. Amount of Variation and Degree of Heterozygosity in Natural Populations of Drosophila Pseudoobscura. *Genetics* 54(2): 595.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Lindner, M. D., A. Vancea, and M. C. Chen. 2016. NIH Peer Review: Scored Review Criteria and Overall Impact. *American Journal of Evaluation* 37(2): 238–249.

Long, P. O. 2011. *Artisan/Practitioners and the Rise of the New Sciences, 1400–1600*. Corvallis, OR: Oregon State University Press.

Longino, H. E. 1990. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton, NJ: Princeton University Press.

Longino, H. E. 2002. *The Fate of Knowledge*. Princeton, NJ: Princeton University Press.

Lyell, C. 1830. *Principles of Geology; Being an Attempt to Explain the Former Changes of the Earth's Surface, by Reference to Causes Now in Operation*. Vol. 1. London: J. Murray.

Macarthur, R. H., and E. O. Wilson. 1963. An Equilibrium Theory of Insular Zoogeography. *Evolution* 17(4): 1–16.

MacArthur, R. H., and E. O. Wilson. 1967. *The Theory of Island Biogeography*. Princeton, NJ: Princeton University Press.

Makdisi, G. 1981. *The Rise of Colleges: Institutions of Learning in Islam and the West*. Edinburgh: Edinburgh University Press.

Maynard Smith, J. 1984. Paleontology at the High Table. *Nature* 309(5967): 401–402.

Melott, A. L., and R. K. Bambach. 2014. Analysis of Periodicity of Extinction Using the 2012 Geological Timescale. *Paleobiology* 40(2): 177–196.

Merton, R. K. 1961. Singletons and Multiples in Scientific Discovery: A Chapter in the Sociology of Science. *Proceedings of the American Philosophical Society* 105(5): 470–486.

Mesoudi, A., K. N. Laland, R. Boyd, B. Buchanan, E. Flynn, R. N. Mccauley, J. Renn, V. Reyes-García, S. Shennan, D. Stout, et al. The Cultural Evolution of Technology and Science. In *Cultural Evolution*, edited by P. J. Richerson and M. H. Christiansen, 194–216. Cambridge, MA: MIT Press.

Mesoudi, A., A. Whiten, and K. N. Laland. 2004. Perspective: Is Human Cultural Evolution Darwinian? Evidence Reviewed from the Perspective of the Origin of Species. *Evolution* 58(1): 1–11.

Miller, B. 2012. When Is Consensus Knowledge Based? Distinguishing Shared Knowledge from Mere Agreement. *Synthese* 190(7): 1293–1316.

Miu, E., N. Gulley, K. N. Laland, and L. Rendell. 2018. Innovation and Cumulative Culture through Tweaks and Leaps in Online Programming Contests. *Nature Communications* 9(1): 2321.

Morrell, J. B. 1972. The Chemist Breeders: The Research Schools of Liebig and Thomas Thomson. *Ambix* 19(1): 1–46.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Mulkay, M. 1976. Problems in the Emergence of New Disciplines. In *Perspectives on the Emergence of Scientific Disciplines*, edited by G. Lemaine, R. Macleod, M. Mulkay, and P. Weingart, 1–23. The Hague: Mouton.

Mulkay, M. J. 1991. *A Sociological Pilgrimage: Studies in the Sociology of Science*. Philadelphia: Open University Press.

Newell, N. 1952. Periodicity in Invertebrate Evolution. *Journal of Paleontology* 26(3): 371–385.

Newell, N., and E. Colbert. 1948. Paleontologist: Biologist or Geologist? *Journal of Paleontology* 22(2): 264–267.

Newell, N. D. 1963. Crises in the History of Life. *Scientific American* 208(2): 76–95.

Newell, N. D. 1967. Revolutions in the History of Life. *Geological Society of America Special Papers* 89: 63–92.

Nosil, P. 2008. Speciation with Gene Flow Could Be Common. *Molecular Ecology* 17(9): 2103–2106.

Nosil, P. 2012. *Ecological Speciation*. New York: Oxford University Press.

Odling-Smee, J. 2007. Niche Inheritance: A Possible Basis for Classifying Multiple Inheritance Systems in Evolution. *Biological Theory* 2(3): 276–289.

Orr, H. A. 1998. The Population Genetics of Adaptation: The Distribution of Factors Fixed during Adaptive Evolution. *Evolution* 52(4): 935–949.

Orr, H. A. 2000. Adaptation and the Cost of Complexity. *Evolution* 54(1): 13–20.

Orr, H. A. 2002. The Population Genetics of Adaptation: The Adaptation of DNA Sequences. *Evolution* 56(7): 1317–1330.

Orr, H. A. 2003. The Distribution of Fitness Effects among Beneficial Mutations. *Genetics* 163(4): 1519–1526.

Orr, H. A. 2005a. The Genetic Theory of Adaptation: A Brief History. *Nature Reviews Genetics* 6(2): 119–127.

Orr, H. A. 2005b. Theories of Adaptation: What They Do and Don't Say. *Genetica* 123: 3–13.

Orr, H. A. 2006. The Distribution of Fitness Effects among Beneficial Mutations in Fisher's Geometric Model of Adaptation. *Journal of Theoretical Biology* 238(2): 279–285.

Orr, H. A. 2018. Evolution, Finance, and the Population Genetics of Relative Wealth. *Journal of Bioeconomics* 20(1): 29–48.

Ospovat, D. 1983. *The Development of Darwin's Theory*. New York: Cambridge University Press.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Pagel, M. 2009. Human Language as a Culturally Transmitted Replicator. *Nature Reviews Genetics* 10(6): 405–415.

Pagel, M., and R. Mace. 2004. The Cultural Wealth of Nations. *Nature* 428(6980): 275–278.

Plutynski, A. 2008. The Rise and Fall of the Adaptive Landscape? *Biology and Philosophy* 23(5): 605–623.

Polanyi, M. 1958. *Personal Knowledge: Towards a Post-Critical Philosophy*. New York: Routledge & Kegan Paul.

Popper, K. 1959. *The Logic of Scientific Discovery*. Abingdon-on-Thames, UK: Taylor & Francis.

Popper, K. R. 1972. *Objective Knowledge: An Evolutionary Approach*. New York: Oxford University Press.

Provine, W. B. 1971. *The Origins of Theoretical Population Genetics*. Chicago: University of Chicago Press.

Raup, D. M. 1966. Geometric Analysis of Shell Coiling: General Problems. *Journal of Paleontology* 40(5): 1178–1190.

Raup, D. M. 1975. Taxonomic Survivorship Curves and Van Valen's Law. *Paleobiology* 1(1): 82–96.

Raup, D. M. 1978. Approaches to the Extinction Problem: Presidential Address to the Society. *Journal of Paleontology* 52(3): 517–523.

Raup, D. M. 1991. *Extinction: Bad Genes or Bad Luck?* New York: W. W. Norton.

Raup, D. M., and S. J. Gould. 1974. Stochastic Simulation and Evolution of Morphology—Towards a Nomothetic Paleontology. *Systematic Zoology* 23(3): 305–322.

Raup, D. M., S. J. Gould, T. Schopf, and D. S. Simberloff. 1973. Stochastic Models of Phylogeny and the Evolution of Diversity. *Journal of Geology* 81(5): 525–542.

Raup, D. M., and J. Sepkoski Jr. 1984. Periodicity of Extinctions in the Geologic Past. *Proceedings of the National Academy of Sciences of the United States of America* 81(3): 801–805.

Reisch, G. A. 2005. *How the Cold War Transformed Philosophy of Science: To the Icy Slopes of Logic*. New York: Cambridge University Press.

Renzi, B. G. and G. Napolitano. 2011. *Evolutionary Analogies: Is the Process of Scientific Change Analogous to the Organic Change?* Newcastle upon Tyne, UK: Cambridge Scholars.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Richards, R. J. 1977. The Natural Selection Model of Conceptual Evolution. *Philosophy of Science* 44(3): 494–501.

Richards, R. J. 1992. *The Meaning of Evolution: The Morphological Construction and Ideological Reconstruction of Darwin's Theory*. Chicago: University of Chicago Press.

Richards, R. J. 2008. Darwin's Theory of Natural Selection and Its Moral Purpose. In *The Cambridge Companion to the Origin of Species*, edited by M. Ruse and R. Richards, 47–66. New York: Cambridge University Press.

Rocke, A. 2001. *Nationalizing Science*. Cambridge, MA: MIT Press.

Rocke, A. J. 1984. *Chemical Atomism in the Nineteenth Century: From Dalton to Cannizzaro*. Columbus: Ohio State University Press.

Rocke, A. J. 2003. Origins and Spread of the "Giessen Model" in University Science. *Ambix* 50(1): 90–115.

Roughgarden, J. 1979. *Theory of Population Genetics and Evolutionary Ecology: An Introduction*. New York: Macmillan.

Rudwick, M. J. S. 1964. The Inference of Function from Structure in Fossils. *British Journal for the Philosophy of Science* 15(57): 27–40.

Schluter, D. 2000. *The Ecology of Adaptive Radiation*. New York: Oxford University Press.

Schopenhauer, A. 1909. *The World as Will and Idea*. Vol. 3. Translated by R. B. Haldane and J. Kemp. London: Kegan Paul, Trench, Trübner & Co.

Schopf, T., D. Raup, and S. J. Gould. 1975. Genomic versus Morphologic Rates of Evolution: Influence of Morphologic Complexity. *Paleobiology* 1(1): 63–70.

Schopf, T. J. M. 1972. Varieties of Paleobiologic Experience. In *Models in Paleobiology*, edited by T. J. M. Schopf, 8–25. San Francisco: W. H. Freeman.

Sepkoski, D. 2012. *Rereading the Fossil Record*. Chicago: University of Chicago Press.

Sepkoski, J. J., Jr. 1978. A Kinetic Model of Phanerozoic Taxonomic Diversity I. Analysis of Marine Orders. *Paleobiology* 4(3): 223–251.

Sepkoski, J. J., Jr. 1979. A Kinetic Model of Phanerozoic Taxonomic Diversity II. Early Phanerozoic Families and Multiple Equilibria. *Paleobiology* 5(3): 222–251.

Sepkoski, J. J., Jr. 1984. A Kinetic Model of Phanerozoic Taxonomic Diversity. III. Post-Paleozoic Families and Mass Extinctions. *Paleobiology* 10(2): 246–267.

Shapin, S. 2010. *Never Pure: Historical Studies of Science as If It Was Produced by People with Bodies, Situated in Time, Space, Culture, and Society, and Struggling for Credibility and Authority*. Baltimore: Johns Hopkins University Press.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Shennan, S. 2013. Long-Term Trajectories of Technological Change. In *Cultural Evolution*, edited by P. J. Richerson and M. H. Christiansen, 143–156. Cambridge, MA: MIT Press.

Sholl, D. A. 1954. Regularities in Growth Curves, Including Rhythms and Allometry. In *Dynamics of Growth Processes*, edited by E. J. Boell, 224–241. Princeton, NJ: Princeton University Press.

Shoval, O., H. Sheftel, G. Shinar, Y. Hart, O. Ramote, A. Mayo, E. Dekel, K. Kavanagh, and U. Alon. 2012. Evolutionary Trade-Offs, Pareto Optimality, and the Geometry of Phenotype Space. *Science* 336(6085): 1157–1160.

Simon, M. 2016. Twitter Nerd-Fight Reveals a Long, Bizarre Scientific Feud. Wired, February 3. https://www.wired.com/2016/02/twitter-nerd-fight-reveals-a-long-bizarre-scientific-feud/.

Simonton, D. K. 1979. Multiple Discovery and Invention: Zeitgeist, Genius, or Chance? *Journal of Personality and Social Psychology* 37(9): 1603–1616.

Simpson, G. G. 1944. *Tempo and Mode in Evolution*. New York: Columbia University Press.

Simpson, G. G. 1945. Tempo and Mode in Evolution. *Transactions of the New York Academy of Sciences* 8: 45–60.

Simpson, G. G. 1949. *The Meaning of Evolution*. New Haven, CT: Yale University Press.

Simpson, G. G. 1953. *The Major Features of Evolution*. New York: Columbia University Press.

Simpson, G. G. 1976. The Compleat Palaeontologist? *Annual Review of Earth and Planetary Sciences* 4(1): 1–14.

Smith, P. H. 2004. *The Body of the Artisan: Art and Experience in the Scientific Revolution*. Chicago: University of Chicago Press.

Stanley, S. M. 1979. *Macroevolution*. San Francisco: Freeman, Cooper and Co.

Stanley, S. M., P. W. Signor III, S. Lidgard, and A. F. Karr. 1981. Natural Clades Differ from "Random" Clades: Simulations and Analyses. *Paleobiology* 7(1): 115–127.

Stauffer, R. C., ed. 1975. *Charles Darwin's Natural Selection: Being the Second Part of His Big Species Book Written from 1856 to 1858*. New York: Cambridge University Press.

Sterelny, K. 2012. *The Evolved Apprentice: How Evolution Made Humans Unique*. Cambridge, MA: MIT Press.

Stroud, J. T., and J. B. Losos. 2016. Ecological Opportunity and Adaptive Radiation. *Annual Review of Ecology, Evolution, and Systematics* 47(1): 507–532.

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Swerdlow, N. 1998. Galileo's Discoveries with the Telescope and Their Evidence for the Copernican Theory. In *The Cambridge Companion to Galileo*, edited by P. Machamer, 244–270. New York: Cambridge University Press.

Tappenden, J. 2005. Proof Style and Understanding in Mathematics I: Visualization, Unification and Axiom Choice. In *Visualization, Explanation and Reasoning Styles in Mathematics*, edited by P. Macouso, K. F. Jørgensen, and S. A. Pedersen, 147–214. Dordrecht: Springer.

Toulmin, S. 1972. *Human Understanding*. Princeton, NJ: Princeton University Press.

Usselman, M., A. Rocke, C. Reinhart, and K. Foulser. 2005. Restaging Liebig: A Study in the Replication of Experiments. *Annals of Science* 62(1): 1–55.

Usselman, M. C. 2003. Liebig's Alkaloid Analyses: The Uncertain Route from Elemental Content to Molecular Formulae. *Ambix* 50(1): 71–89.

Van Valen, L. 1973. A New Evolutionary Law. *Evolutionary Theory* 1(1): 1–30.

Wallace, B. 1991. *Fifty Years of Genetic Load: An Odyssey*. Ithaca, NY: Cornell University Press.

Westfall, R. S. 1980. *Never at Rest: A Biography of Isaac Newton*. New York: Cambridge University Press.

Westman, Robert S. 2011. *The Copernican Question: Prognostication, Skepticism, and Celestial Order*. Berkeley: University of California Press.

White, J. F., and S. J. Gould. 1965. Interpretation of the Coefficient in the Allometric Equation. *American Naturalist* 99(904): 5–18.

Wray, K. B. 2011. *Kuhn's Evolutionary Social Epistemology*. New York: Cambridge University Press.

Zyla, P. A., R. M. Barnett, J. Beringer, O. Dahl, D. A. Dwyer, D. E. Groom, C.-J. Lin, K. S. Lugovsky, E. Pianori, D. J. Robinson, et al. 2020. Review of Particle Physics. *Progress of Theoretical and Experimental Physics* 2020(8).

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Index

Adaptation, 37, 71, 81, 88, 96–101, 105–107, 137, 175, 194, 226, 230, 236
Adaptive walk, 108, 113–134, 139–140, 303
Adaptive zone, 182–184, 189, 218, 286
Allometric equations, 228–229, 231, 233–235, 253
Allopatric model, 160–161, 238–239, 241–246
Anagenesis, 248
Atomic theory, 56, 65, 147
Attractors, 107, 136

Beatty, John, 226
Berzelius, J. J., 59, 65–66
Biodiversity, 272
Block, Ned, 203
Bolnick, Daniel, 157, 160, 167
Brahe, Tycho, 180
Branching, 30, 140–141, 143–145, 147–149, 151–156, 159, 161–163, 165, 167, 169–171, 173, 175, 177–178, 197, 199, 200–201, 203–204, 209, 241, 261, 264, 266, 274, 276, 294. *See also* Speciation
Brandon, Robert, 45, 178
Breeding, 12–13, 55–56, 61, 78, 157, 161, 238
Brush, Stephen, 251
Burkhardt, Richard, 274

Cambrian period, 197, 294, 296
Campbell, Donald, 20–21, 23–25
Carnap, Rudolf, 166
Catastrophe, 177, 181, 184–185, 275, 284, 290
Causality, 33, 38
Chicxulub, 180–181, 284
Child-rearing advice, 55
Clades, 151–152, 154–156, 185–186, 259, 262–266, 268–269, 293, 307–308
Cloud, Preston, 219
Colbert, Edwin, 213–215
Combustion, 56, 58–59, 133
Conceptual change, 24–25, 89
Consensus practice, 49, 63, 186, 194
Content bias, 59
Contingency, 79, 84, 223–226, 249, 258, 269
Convergence, 95, 108–110, 112–113, 127, 133, 135, 269
Convergent evolution, 106
Copernican model, 132, 186
Cope's Law, 231
Cretaceous period, 180–181, 185, 189, 271, 284, 292
Crises, Kuhnian, 87, 177, 185–187, 191, 194–195

Daltonian atomism, 65, 133, 147
Darwin, Charles, 37, 105–106, 114, 131, 134–135, 154, 236, 272–283, 301

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Darwinian evolution, 11, 15, 139, 183
Databases, 219
Demes, 24
Dimopoulos, Savas, 130–131
Dinosaurs, 180–181, 185, 271, 284
Directionality, 29–37, 45–46, 49, 64, 66, 68–69, 71–72, 76–82, 86, 90, 96, 101, 106, 108, 113, 135, 139, 158, 215, 263
Directional refinement, 33, 46–47, 49–50, 55, 64–66, 73–74, 76, 80, 90–91, 108, 113, 139–141, 173–175
Discrete heritable unit, 43
Divergence, Principle of, 154, 157–161, 197–198, 200, 222, 277–278
Dobzhansky, Theodosius, 202–203, 214
Dollo, Louis, 223
Dollo's Law, 223–225, 253–254, 260, 308
Donoghue, Michael J., 149–150
Drift, 72–73, 96–97, 122, 161–162, 177, 200, 218
Dumas, Jean-Baptiste, 62, 66

Ecological inheritance, 76
Ecosystems, 50, 181–184, 189, 203–204
Eldredge, Niles, 213, 237–250, 266
Endemism, 75–76
Epistemic cultures, 47
Erwin, Doug, 181, 183–184
Evolutionary paleontology, 211–212, 232, 268
Evolution of languages, 43
Evolvability, 118–121
Exemplar, 56, 197
Extinction
  Darwin's theory of, 280, 283, 285, 304
  mass, 178–181, 186, 188–189, 200, 219, 271–272, 291–293, 298–300

Feynman, Richard, 95
"Field of bullets," 285
Fisher, R. A., 5, 114–115, 123–124, 127, 129, 214, 246
Fisherian runaway, 79
Fitzpatrick, Daniel, 157, 160, 167
Foothold, 130–133, 140, 188
Forman, Paul, 33–34
Fossil, 223, 228–229, 231–232, 245, 251–253, 265, 270, 275
  record, 150, 184, 214–219, 222, 224–225, 230, 237–239, 241, 243, 246–250, 252, 266–268, 272, 279–280, 282, 285–286, 288, 296–298, 301–305, 307
  taphonomic, 247
Fruitfulness, 43, 82, 86, 89, 124, 127, 136, 192, 225–228, 230–231, 234–235, 253, 268–269
Futuyma, Douglas, 246, 248, 275, 280

Galileo, 130, 132, 135
Galison, Peter, 1
Gambler's Ruin Model, 286–289, 292
Gay-Lussac, Joseph Louis, 56–57, 59
Genus (genera), 92, 180–181, 184, 189, 219, 273, 275–279, 287–288, 291–292, 298–299
Geometry, 54, 115, 118, 123, 127, 129
Giessen, 56, 62
Gillespie, Charles, 115
Godfrey-Smith, Peter, 11–13, 15–16, 20–21, 37–38, 41, 73, 107, 139, 193
Goodman, Nelson, 12
Gould, Stephen Jay, 29, 211–214, 219–226, 228–237, 242–253, 258–260, 262–265, 269, 271, 284, 294, 298, 303–304, 306–308
Gradualism, 218, 222–223, 237–239, 241, 243
Grosholz, Emily, 39

Hacking, Ian, 1, 139, 147
High-fidelity copying, 45–49, 53, 56, 60, 64–67, 69, 71, 109
Holmes, Larry, 57, 60–64, 67
Holton, Gerald, 95, 116, 126, 170–172

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Homophony, 222
Hooker, Joseph, 124
Hull, David, 14–25, 24–26, 39, 43
Huss, John, 267

Incommensurability, 17–19, 139, 198–199, 201–204
Inductivism, 223, 225, 230, 249–250, 254
Interactors, 24, 26
Intrinsicality, 38, 41, 73, 99, 191–193, 196, 217, 271, 303
Iridium, 284
Islam, 46
Isolation, 22, 71–73, 78–85, 91, 94, 96, 103, 167, 172–173, 200, 216–217, 238, 248
Isomerism, 65–66, 68

Jablonski, David, 181, 184–186, 291–294, 304–305
Janiak, Andrew, 189

Kaliappart, 58–60, 63–65
Kepler, Johannes, 31
Kimura, Motoo, 114–115
Kitcher, Philip, 39, 49, 124, 144
Knight, J. Brookes, 215–216
Knorr Cetina, Karin, 47
Knowledge production, 20–21, 24–25, 28, 33–34, 36, 42, 61–63, 69, 82–83, 85–87, 91, 100–104, 108, 127, 133–134, 154, 156–157, 163, 168, 171, 185–187
Kohn, David, 274, 276–279
Kuhn, Thomas, 17–19, 22–24, 28, 30, 36, 46–47, 54–55, 63, 81–92, 108, 111, 120–129, 137–140, 144–145, 164–165, 172–175, 177–179, 185–199, 201–205

Lagrange, Joseph-Louis, 130
Laland, Kevin, 88
Lamarck, Jean-Baptiste, 274–275
Lavosier, Antoine-Laurent de, 120, 135, 186
Lawn mowing, 38, 40
Lewontin, Richard, 202, 226, 236, 302
Lexical change, 18, 82–83, 199–200
Lexical diversity, 82–83, 173, 199–203
Liebig, Justus, 48, 56–69
Lineages, 153–156, 163, 183–184, 228, 238, 241–242, 256, 261–262, 264–266, 279–281, 288, 293
Lyell, Charles, 220–221, 273–275

MacArthur, Robert, 260–262, 273, 294
Mace, Ruth, 167, 217
Macroevolution, 156, 263, 304–305
Makdisi, George, 46
Malthus, Thomas Robert, 131, 274
Marine Biological Laboratory (MBL), 258–270, 284–286, 294, 303, 307
Markov
  chains, 224, 253
  process, 258
Mate preference, 79–80, 85–86, 160, 233–234
Maynard Smith, John, 247–248, 271–272, 284, 300–302
Mayr, Ernst, 202
Mechanics, 31, 33–34, 54, 130, 187–189, 229, 252
Merton, Robert, 109–112, 127, 135–136
Mesoudi, Alex, 40, 42, 50, 59
Mesozoic era, 184
Migration, 73, 78, 96–97, 170–172, 254, 260–261, 273, 283
Modern Synthesis, 214, 216, 222, 228, 247, 251
  Modern Synthetic picture, 114
  Modern Synthetic theory, 84
Monopoly, 182, 217, 284, 305–306
Morphospace, 37, 45–46, 106–107, 121, 135, 158, 183, 187, 193, 222–223, 232, 245, 253, 269, 276

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Morrell, Jack, 56, 59–60
Mulkay, Michael, 49, 102, 127, 169–171

Newell, Norman, 177, 213–216, 219–223, 237
Newton, Isaac, 31, 122–123, 135–136, 143, 165, 187–189, 199–201, 236
Niche construction, 14, 72–93, 97–103, 164–166
Nomothetic undertones, 223, 230, 253, 307
Normal science, 18, 54–55, 63, 82, 85, 125, 128, 137, 139, 175–176, 179, 187, 190–196, 198
Nosil, Patrik, 84
Novelty, 52, 88, 91–93, 179, 184, 188, 195–197, 199, 287

On the Origin of Species (Darwin), 13, 37, 105, 124, 131, 134, 163, 176, 213, 272, 277–280, 283, 289, 301–302
Optimality, 129, 229, 232–236
Organic chemistry, 56–64, 66, 68
Orr, H. Allen, 12, 115–118, 120, 123–124, 126–129, 133, 140, 303
Orthogenesis, 251
Ospovat, Dov, 84, 131, 274–276

Pagel, Mark, 43, 167, 217
Paleobiology, 170, 213, 237, 242, 254–255, 264, 271
Paleoecology, 219
Paleontology, 40, 284–285, 294, 300–308
Paleozoic era, 184, 239, 296
Paradigms, 18–19, 38–40, 46, 53, 55, 63, 71, 82, 85–86, 96, 103, 108, 120–125, 127–129, 139, 165–167, 173–176, 178–179, 186–189, 191–192, 194, 197, 234
Paradigm shifts, 18–19, 40, 63, 89, 194–198
Permian period, 181, 184
Phacops rana, 239, 241
Phanerozoic, 293–295, 296
Phyletic
  evolution, 218, 244
  gradualism, 222, 237–244
  model of transformation, 238, 239–241
Phylogenetic
  accidents, 232
  results, 224, 225
  tools, 43
  trees, 238, 263
Popper, Karl, 15, 20–21, 23–24
Population genetics, 247–248, 251–252, 268–272, 304–305
Populations/communities, 14, 22, 36–38, 41, 46–47, 50, 55, 69, 71–75, 77, 90–91, 93–94, 100–101, 107, 111–112, 117, 119, 121, 127, 134, 141, 154–155, 162, 173, 178, 194–195, 203, 209, 211, 217, 227, 254, 272
Power function, 228–235, 253
Practice space, 108–113, 123–124, 128–130, 136, 146, 168, 187–189
Progress
  horizontal, 30, 36, 58, 140, 146–148, 154–156, 191, 279
  vertical, 29, 36, 47–48, 55, 139–140, 146–148, 155, 167, 191
Pseudodirectionality, 30–31
"Pummeling of the Predecessors," 13, 26
Punctuated equilibrium, 237, 241–250, 253–254, 266–267, 275, 307
Puzzles, 15, 28, 85–86, 134, 201, 252, 307

Quantum evolution, 33–34, 147, 187, 217–218, 246

Radiation, adaptive, 261, 298–301
Randomness, 20–21, 34, 48, 122, 159–161, 170, 191, 218, 257–269, 286, 296, 298

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Raup, David, 98, 149, 229, 232, 236, 243, 253, 259, 261–263, 272–273, 284–286, 292, 294, 298–300, 303–304
Replicators, 24, 26
Research funding, 52, 88–90, 95, 99, 164
Research problems, 50–53, 55–63, 72, 85–88, 92–99, 102, 120, 124–125, 131–132, 167–173, 185–192, 210, 227–228, 303, 306
Richards, Robert J., 20
Rocke, Alan, 59, 62, 65, 67–68, 133
Rudwick, Martin, 229–230, 234–235

Saltationist model, 241, 251
Sampling error, 97–98, 161–162
Schindewolf, Otto, 220, 222
Schopf, Thomas, 242–243, 249, 254–256, 258–260, 284, 286, 302–303
Science of form, 228–237, 249
Scientific education, 49, 53–54, 88, 96, 194
Scientific genius, 31, 135–136
Scientific practice, 39, 41, 43, 47–48, 50, 55–56, 66, 69, 71–72, 80–81, 90, 96, 107–108, 113, 127, 134, 139, 153, 172–174, 190–191, 194, 198, 200–201, 204, 209
Scientific progress, 28, 30, 34, 36, 39, 40, 47, 62, 72, 81–83, 108, 123, 146, 175, 196, 220, 249, 259
Scientific Revolution, 199
Scientific revolutions, 18, 31, 175, 176, 189, 195, 199, 205
Selection, 26, 37, 41, 45–46, 48, 52, 72, 97–99, 105–107, 114, 134, 156–157, 169, 173, 229, 267, 289
Sepkoski, David, 213, 219, 254, 258, 263, 285–286, 308
Sepkoski, J. J., Jr., 293–298
Shapin, Steven, 3–4
Significance, 94–95, 100, 156, 163
Simberloff, Daniel, 259
Simonton, D. Keith, 110–111
Simpson, G. G., 149–150, 182, 202, 214–218, 225, 230, 238, 246–247, 251, 254
Simulations, 262–264, 267–268, 293–294
Sociobiology, 226, 233–234
Sociobiology Wars, 233–234
Specialization, 162–163, 167–168, 170–173, 209–210
Speciation, 17, 81, 84, 155, 157, 160–163, 167, 177, 181, 199, 200, 210, 216, 238–249, 261, 272, 278, 287, 295, 312n9. *See also* Branching
Stability, 144–147, 151
Stanley, Steven, 264–267
Stasis, 145, 243–250, 253, 266, 280
Stella nova of 1572, 180
Stratigraphy, 213, 249, 259, 270, 284
Structural chemistry, 64, 66, 68
Structure of Evolutionary Theory, The (Gould), 223, 228, 232
Structure of Scientific Revolutions, The (Kuhn), 1, 17–24, 30, 47, 53, 63, 81–86, 92, 137, 178, 195, 197–199
Survival, 76, 98, 152, 184, 216, 292
Survivorship, 291–293
Systematics, 25–26

Tappenden, James, 81, 131
Taxonomy, 18, 25, 125, 149–150, 182–184, 189, 199, 216–219, 230, 241–243, 256–257, 262–269, 276–277, 283, 286, 288–289, 293–294, 296, 301–302, 308
Tempo and Mode (Simpson), 214–215, 217–218
Toulmin, Stephen, 15, 22–24, 39, 43

Uniformitarianism, 220, 259
  methodological, 220–221
  substantive, 220–222
Usselman, Mel, 56, 58

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023

Valentine, James, 183–184, 303
Van Valen, Leigh, 256–259, 286
Veterans and apprentices, 39, 46–47, 49, 53, 55, 272

Ward, Pete, 150–151
Westfall, Richard S., 136, 143
Williston's Law, 231
Wilson, Edward O., 254, 259–262, 273, 294
Wöhler, Friedrich, 65–66
Wright, Sewall, 202, 214, 218, 245

Downloaded from http://direct.mit.edu/books/book-pdf/2086151/book_9780262371599.pdf by guest on 06 August 2023